

**In The United States Patent and Trademark Office  
On Appeal From The Examiner To The Board  
of Patent Appeals and Interferences**

In re Application of: Craig H. Rowland  
Application No.: 10/685,726  
Filed: October 15, 2003  
Confirmation No.: 5392  
Art Unit: 2131  
Examiner: Aravind K. Moorthy  
Title: Method and System for Reducing the False Alarm Rate of  
Network Intrusion Detection Systems

**Mail Stop: Appeal Brief - Patents**  
Commissioner for Patents  
P.O. Box 1450  
Alexandria, Virginia 22313-1450

Dear Sir:

**Appeal Brief**

Appellant has appealed to the Board of Patent Appeals and Interferences (the "Board") from the decision of the Examiner mailed May 30, 2007, finally rejecting all pending Claims 1-21. An Advisory Action was mailed August 15, 2007. Appellant filed a Notice of Appeal and a Pre-Appeal Brief Request for Review, with the then-applicable statutory fee of \$500.00, on August 30, 2007. A Notice of Panel Decision from Pre-Appeal Review was mailed November 5, 2007, indicating that the case should proceed to the Board because there is at least one actual issue for appeal. Appellant respectfully submits this Appeal Brief with the statutory fee of \$510.00.

**Real Party in Interest**

This Application is currently owned by Cisco Technology, Inc. as indicated by an assignment recorded on October 15, 2003 in the Assignment Records of the United States Patent and Trademark Office (PTO) at Reel 014628, Frame 0166 (3 pages).

**Related Appeals and Interferences**

To the knowledge of Appellant's counsel, there are no appeals, interferences, or judicial proceedings that are related to or will directly affect, be directly affected by, or have a bearing on the Board's decision regarding this Appeal.

**Status of Claims**

Claims 1-21 are pending in this Application and stand rejected pursuant to a Final Office Action dated May 30, 2007 ("Final Office Action"), and are all presented for appeal. All pending claims are shown in Appendix A, along with an indication of the status of those claims.

**Status of Amendments**

All amendments submitted by Appellant have been entered by the Examiner prior to the mailing of the Final Office Action.

**Summary of Claimed Subject Matter**

In certain embodiments, the present invention provides a system 100 (*see, e.g.,* FIGURE 1) for reducing the false alarm rate of a network intrusion detection system (“NIDS”) 108 by utilizing an offline passive analysis tool 110. In the illustrated embodiment, system 100 includes NIDS 108 coupled to a link 106 that communicatively couples an unprotected network 102 with a protected network 104, a network 120 that couples NIDS 108 with passive analysis tool 110, a dynamic host configuration protocol (“DHCP”) server 122 coupled to passive analysis tool 110, and a network administrator 112 that utilizes passive analysis tool 110. (Spec. at 6:6-18))

Unprotected network 102 may be any suitable network external to protected network 104. An example of unprotected network 102 is the Internet. Protected network 104 may be any suitable network, such as a local area network, wide area network, virtual private network, or any other suitable network desired to be secure from unprotected network 102. Link 106 couples unprotected network 102 to protected network 104 and may be any suitable communications link or channel. Communications link 106 may be operable to transmit data in “packets” or another suitable form between unprotected network 102 and protected network 104; however, communications link 106. (Spec. at 6:19-7:2)

NIDS 108 may be any suitable network-based intrusion detection system operable to analyze data packets transmitted over communications link 106 in order to detect any potential attacks on protected network 104. Generally, network intrusion detection systems include one or more sensors having the ability to monitor any suitable type of network having any suitable data link protocol. In addition, some network intrusion detection systems are passive observers of network traffic and do not have their own network address. (Spec. at 7:3-14)

In certain embodiments, sensors associated with NIDS 108 are operable to examine data packets on an Internet Protocol (IP) network using any suitable protocol, such as Transmission Controlled Protocol (TCP), User Datagram Protocol (UDP), and Internet Controlled Message Protocol (ICMP). Upon detection of a possible attack on protected network 104, NIDS 108 is operable to generate an alarm indicating that an attack on protected network 104 may have occurred. Alarm trigger packets are then transmitted to passive analysis tool 110 over network 120 along with one or more other data packets associated with the alarm for analysis. (Spec. at 7:15-27)

In certain embodiments, passive analysis tool 110 is a backend application that receives, via network 120, one or more data packets from NIDS 108 and, using the information associated with the data packets, determines if an attack is real or merely a false alarm. These data packets, which may be any suitable portion of an information stream, include characteristics of the alarm, such as an attack type and an operating system (OS) fingerprint for the target host so that passive analysis tool 110 may analyze the potential attack without having access to the network stream on link 106. (Spec. at 7:28-8:8)

In this manner, passive analysis tool 110 significantly lowers the false alarm rate for NIDSs, such as NIDS 108, in the network environment and lowers the requirement of personnel, such as network administrator 112, monitoring these systems to respond to every alarm. In addition, passive analysis tool 110 may reside anywhere in an enterprise and may be used with different types of NIDS, even legacy NIDS that do not support passive OS fingerprinting. Passive analysis tool 110 may also facilitate the analysis of target hosts that are behind strong or impenetrable firewalls. (Spec. at 8:9-20)

In the illustrated embodiment, passive analysis tool 110 (described in greater detail with respect to FIGURE 2) is coupled to NIDS 108 via network 120, which may be any suitable network, or combination of networks, such as a local area network, wide area network, global network, virtual private network, or any other suitable network. Network administrator 112 may be any suitable personnel that utilizes passive analysis tool 110 in order to monitor potential attacks on protected network 104 and respond thereto, if appropriate. As an example, network administrator 112 may have passive analysis tool 110 residing on his or her computer in order to receive filtered alarms from passive analysis tool, as denoted by reference numeral 114. (Spec. at 8:21-9:4)

Some embodiments of the invention provide numerous technical advantages; other embodiments may realize some, none, or all of these advantages. For example, the false alarm rate of NIDS may be substantially reduced or eliminated, which may lead to a lower requirement of personnel monitoring of NIDS to respond to alarms. This may be facilitated by a system in which there is no need to access the network stream to determine the operating system type of the target host. The system may reside anywhere in an enterprise and may be used with different types of NIDS, even legacy NIDS sensors that do not support passive OS fingerprinting. Such a system may free up the NIDS so that it runs more efficiently and at a

faster speed. In addition, an offline passive analysis system may facilitate the analysis of target hosts that are behind strong or impenetrable firewalls. (Spec. at 3:16-4:2)

FIGURE 2 is a block diagram illustrating various example functional components of passive analysis tool 110. In the illustrated embodiment, passive analysis tool 110 includes an alarm input layer 202, an alarm interpretation layer 204, a target cache look-up 206, a passive offline fingerprinting mechanism 208, and an alarm output layer 210. (Spec. at 9:5-17)

Alarm input layer 202 is generally responsible for receiving the data packets from NIDS 108 and determining if the alarm format is valid. If the alarm format is invalid, then the alarm is disregarded. If the alarm format is valid, then the alarm is sent to alarm interpretation layer 204. Alarm input layer 202 is preferably designed to be NIDS vendor independent so that it may accept alarms from multiple NIDS sources concurrently with no modification. Alarm input layer 202, in one embodiment, may also accept alarms from legacy NIDS that do not support passive OS fingerprinting. (Spec. at 9:18-29)

Generally, alarm interpretation layer 204 receives the data packets from alarm input layer 202 and performs an analysis on the alarm. In one embodiment, alarm interpretation layer 204 determines whether the alarm is from a supported NIDS vendor. If the alarm is not from a supported NIDS vendor, an alert is generated and the alarm is disregarded. If the alarm is from a supported NIDS vendor, then alarm interpretation layer 204 is responsible for identifying the attack type, relevant operating system type being attacked (e.g., Microsoft Windows, Sun Solaris, Linux, UNIX, etc.), the source address, target network address, the alarm severity, the alarm description, and any other suitable parameters associated with the alarm. Some of this information is used by passive analysis 110 to test if the alarm is real or false. (Spec. at 9:30-10:15)

Target cache look-up 206 indicates that a look-up is performed by passive analysis tool 110 in order to determine if the target host has already been checked for the particular attack indicated by the alarm. The look-up may be performed in any suitable storage location, such as a local state table or database. (Spec. at 10:16-21)

Passive offline fingerprinting mechanism 208 performs a passive analysis of the target host by identifying, from the received data packets, the operating system fingerprint of the target host, which includes the operating system type, and comparing the operating system type to the attack type. An advantage of this type of OS fingerprinting is that it requires no



access to the network stream. Passive offline fingerprinting mechanism 208 may store this information in a suitable storage location. (Spec. at 10:22-11:2)

Alarm output layer 210 is responsible for taking the analyzed data from passive analysis tool 110 and either escalating or de-escalating the alarm. In other words, alarm output layer 210 functions to report a valid alarm; i.e., that a particular target host is vulnerable to an attack. A valid alarm may be reported in any suitable manner, such as a graphical user interface, a log file, storing in a database, or any other suitable output. In one embodiment, a valid alarm is automatically reported to network administrator 112 via any suitable method. (Spec. at 11:3-12)

FIGURE 3 is a flow chart illustrating an example method for reducing the false alarm rate of network intrusion detection systems according to one embodiment of the present invention. The example method begins at step 300 where one or more data packets associated with an alarm is received from NIDS 108 by passive analysis tool 110. These data packets may be any suitable portion of an information stream and may be communicated to passive analysis tool 110 via network 120 or other suitable communication means. From the data packets, passive analysis tool 110 identifies the attack type, as denoted by step 302, and an operating system fingerprint of the target host, as denoted by step 304. The operating system type of the target host may be identified by passive analysis tool 110 from the OS fingerprint, as denoted by step 306. (Spec. at 11:17-12:2)

The attack type and the operating system type of the target host are compared at step 308 by passive analysis tool 110. At decisional step 310, it is determined whether the operating system type of the target host matches the attack type. If there is a match, then a confirmed alarm is reported by step 312. In one embodiment, the confirmed alarm is automatically reported to network administrator 112 in any suitable manner. If there is no match, then a false alarm is indicated, as denoted by step 314. For example, if the attack type is for a Windows system and the operating system fingerprint shows a Windows host, then the alarm is confirmed. However, if the attack type is for a Windows system and the operating system fingerprint shows a UNIX host, then this indicates a false alarm. (Spec. at 12:3-18)

Although the method outlined in FIGURE 3 is described with reference to passive analysis tool 110 comparing an operating system type with an attack type, other suitable characteristics of the operating system may be compared to relevant characteristics of the

attack type in order to determine if the alarm is real or false. This depends on the type of information passed from NIDS 108 via the data packets. (Spec. at 12:19-26)

Thus, passive analysis tool 110 is intelligent filtering technology that screens out potential false alarms while not requiring access to protected network 104. Alarm inputs are received from a deployed NIDS, such as NIDS 108, and analyzed to determine if an attack is real or a false alarm. (Spec. at 12:27-13:2)

FIGURE 4 is a flowchart illustrating an example method that may be used in conjunction with the example method outlined in FIGURE 3 in certain embodiments of the present invention. At step 400, DHCP server 122 (FIGURE 1) is monitored by passive analysis tool 110. The present invention contemplates any suitable dynamic configuration protocol server being monitored by passive analysis tool 110. At step 402, lease activity is detected by passive analysis tool 110. At decisional step 404 it is determined whether a lease issue is detected or a lease expire is detected. (Spec. at 13:3-14)

If a lease expire is detected by passive analysis tool 110, then the system cache is accessed, as denoted by step 406. At step 408, it is determined whether the target address associated with the lease expire is found in the system cache. If the target address is found in the system cache, then the entry is purged, at step 410, from the system cache. Passive analysis tool 110 then continues to monitor the DHCP server. If a target address is not found in the system cache, then the lease expire is disregarded, as denoted by step 412. Passive analysis tool 110 continues to monitor the DHCP server. (Spec. at 13:15-26)

Referring back to step 404, if a lease issue has been detected, then the system cache is accessed, as denoted by step 414. At step 416, it is determined whether the target address associated with the lease issue is found in the system cache. If the target address is found, then the entry is purged, at step 418. If the target address is not found in the system cache, then passive analysis tool 110 continues to monitor the DHCP server. (Spec. at 13:27-14:4)

The method outlined in FIGURE 4 addresses the dynamic addition, subtraction, or modifying of hosts in protected network 104 so that prior knowledge of protected network 104 is not required. This saves considerable time and money and is more accurate than prior systems in which prior knowledge of the network is required. Passive analysis tool 110 may more accurately keep track of changes regarding the target hosts of protected network 104. (Spec. at 14:5-13)

With regard to the independent claims currently under Appeal, Appellant provides the following concise explanation of the subject matter recited in the claim elements. For brevity, Appellant does not necessarily identify every portion of the Specification and drawings relevant to the recited claim elements. Additionally, this explanation should not be used to limit Appellant's claims but is intended to assist the Board in considering the Appeal of this Application.

For example, independent Claim 1 recites the following:

A computerized method for reducing the false alarm rate of network intrusion detection systems, comprising:

receiving, from a network intrusion detection sensor, one or more data packets associated with an alarm indicative of a potential attack on a target host (e.g., Figures 1-3; Spec. at 7:3-8:8, 9:5-29, 11:17-26);

identifying characteristics of the alarm from the data packets, including at least an attack type and an operating system fingerprint of the target host (e.g., Figures 1-3; Spec. at 7:28-8:20, 9:5-10:15, 11:17-12:2);

identifying the operating system type from the operating system fingerprint (e.g., Figures 1-3; Spec. at 7:28-8:8, 9:30-11:2, 11:17-12:2);

comparing the attack type to the operating system type (e.g., Figures 1-3; Spec. at 7:28-8:8, 10:22-11:2, 12:3-18); and

indicating whether the target host is vulnerable to the attack based on the comparison (e.g., Figures 1-3; Spec. at 7:28-8:8, 11:3-10:12, 12:3-18).

The citations listed above with respect to independent Claim 1 are also applicable to independent Claims 7 and 16, which are directed to systems (Claim 16 being in means-plus-function form).

**Ground of Rejection to be Reviewed on Appeal**

Are Claims 1-21 patentable under 35 U.S.C. § 102(e) over U.S. Patent 7,152,105 to McClure, et al. ("*McClure*")?

### **Argument**

For at least the following reasons, the Examiner's rejections of Claims 1-21 are improper and should be reversed by the Board.

#### **I. Overview**

Claims 1-21 stand rejected under 35 U.S.C. § 102(e) as being unpatentable over *McClure*, a copy of which is attached as Appendix B. Appellant respectfully submits that at least the portions of *McClure* cited by the Examiner fail to disclose, teach, or suggest each and every limitation recited in Appellant's claims. Appellant respectfully submits that these rejections are therefore improper and should be reversed by the Board.

#### **II. Standard for Demonstrating Anticipation**

"A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference." *Verdegaal Bros. v. Union Oil Co. of California*, 2 U.S.P.Q.2d 1051, 1053 (Fed. Cir. 1987); M.P.E.P. ch. 2131. In addition, "[t]he elements must be arranged as required by the claim." *Richardson v. Suzuki Motor Co.*, 9 U.S.P.Q.2d 1913, 1920 (Fed. Cir. 1989); *In re Bond*, 15 U.S.P.Q.2d 1566 (Fed. Cir. 1990); M.P.E.P. ch. 2131.

#### **III. The Claims are Allowable over *McClure***

##### **A. Independent Claims 1, 15, and 29 (and Their Respective Dependent Claims) are Allowable Over *McClure***

Claim 1, which Appellant discusses as an example, recites:

A computerized method for reducing the false alarm rate of network intrusion detection systems, comprising:

receiving, from a network intrusion detection sensor, one or more data packets associated with an alarm indicative of a potential attack on a target host;

identifying characteristics of the alarm from the data packets, including at least an attack type and an operating system fingerprint of the target host;

identifying the operating system type from the operating system fingerprint;

comparing the attack type to the operating system type; and

indicating whether the target host is vulnerable to the attack based on the comparison.

At a minimum, Appellant respectfully submits that the cited portions of *McClure* do not disclose, teach, or suggest the following limitations recited in Claim 1:

- receiving, from a network intrusion detection sensor, one or more data packets associated with an alarm indicative of a potential attack on a target host;
- identifying characteristics of the alarm from the data packets, including at least an attack type and an operating system fingerprint of the target host;
- comparing the attack type to the operating system type; and
- indicating whether the target host is vulnerable to the attack based on the comparison.

For example, the cited portions of *McClure* fail to disclose, teach, or suggest “receiving, from a network intrusion detection sensor, one or more data packets associated with an alarm indicative of a potential attack on a target host,” as recited in Claim 1. The Examiner relies on the passages at Col. 24, Lines 11-21 (“Passage A”), Col. 17, Line 29 - Col. 18, Line 50 (“Passage B”), and Col. 31, lines 19-36 (“Passage C”) of *McClure* as allegedly disclosing these limitations. (See Final Office Action at 3; Advisory Action, Continuation Sheet; and Notice of Panel Decision from Pre-Appeal Brief Review, Attachment)

According to Passage A:

Sometimes, in order to “force” a response from the target computer, *an intruder* may send a **malformed packet** to a target port. While this known technique increases the likelihood that an open UDP port on the target computer can be identified, this technique also substantially increases the likelihood that the malformed packet could **damage the target computer**. Also, firewalls or routers may detect and filter out malformed packets, and such packets can alert the target network of an attempted security breach.

***The intelligent UDP port scanning test in accordance with this embodiment of the present invention employs an efficient, less intrusive and more accurate method for scanning UDP ports on a target computer.***

(*McClure* at 24:11-26 (emphasis added).) Passage A relates to a technique for discovering host computers (live target computers), particularly to a technique for applying an Intelligent UDP Port Scanning test to each IP address on a scan list. (*McClure* at 22:31-38, 23:54, and 24:21-27)

Passage B discloses packets used to identify an operating system. (*McClure* at 17:36-18:3; *see also McClure* at 18:43-44 (stating “[b]elow is an example exchange of packets when performing an OS [operating system] identification”).)

Passages A and B fail to disclose, teach, or suggest “receiving, from a network intrusion detection sensor, one or more data packets associated with an alarm indicative of a potential attack on a target host,” as recited in Claim 1.

First, as explained in Passage A, a malformed packet is apparently sent by *an intruder*. *An intruder* sending a malformed packet, as is relied upon by the Examiner to allegedly disclose the one or more packets recited in Claim 1 is different than “receiving, *from a network intrusion detection sensor*, one or more data packets associated with an alarm indicative of a potential attack on a target host” of Claim 1 (emphasis added).

Second, the packets of Passage A are completely different from the packets of Passage B, and thus cannot disclose, teach, or suggest the packets of Claim 1. Passage A clearly discloses that a technique of forcing a response using a malformed packet is not used in the system of *McClure* because the technique could damage a computer. That is, the malformed packets of Passage A are not used in the system described by Passage B.

Furthermore, *McClure* discloses that the packets of Passage B are RFC-compliant TCP packets. (*McClure* at 14:41-56; *see also McClure* at 16:57-17:4.) The RFC-compliant TCP packets of Passage B, however, are not the malformed packets of Passage A:

The use of RFC-compliant TCP packets advantageously *reduces the probability that the detection packets are blocked by a router or firewall, and greatly reduces the probability that the detection packets will cause damage or crashes at the target computer.*

(*McClure* at 16:62-67 (emphasis added).) That is, the packets of Passage B greatly reduce the problems associated with the malformed packets of Passage A. As a result, Passages A and B of *McClure* fail to disclose “receiving, from a network intrusion detection sensor, one

or more data packets associated with an alarm indicative of a potential attack on a target host” of Claim 1.

According to Passage C:

In the decision step 730, the process determines whether all the live target computers have been processed in TCP full connect scanning or whether all the batches of live target computers have been processed in TCP SYN scanning. If all the target computers or all the batches of target computers have been processed, the process ends. Otherwise, the process proceeds to a TCP service scan routine 740 wherein the process uses a TCP service discovery list 742 to identify the TCP service ports to be examined for each target computer. As described above, TCP packets are sent to the identified TCP service ports of each target computer, and the target computer vulnerability database 714 is updated for each target computer in accordance with whether a response is received or is not received from each target computer for each TCP service port scanned and using the known vulnerability database to obtain the vulnerability information for the particular TCP service ports that are determined to be open.

(*McClure* at 31:19-36)

Therefore, Passage C of *McClure* also fails to disclose “receiving, from a network intrusion detection sensor, one or more data packets associated with an alarm indicative of a potential attack on a target host” of Claim 1. Nowhere does the cited portion disclose, teach or suggest receiving any message from a network intrusion detection sensor, let alone receiving “one or more data packets associated with an alarm indicative of a potential attack on a target host,” as recited in Claim 1. The Examiner states that Passage C discloses that “[p]ackets are received indicative of vulnerabilities.” (Notice of Panel Decision from Pre-Appeal Brief Review, Attachment) It appears to Applicant, however, that Passage C states that “TCP packets are sent to the identified TCP service ports [identified using TCP discovery list 742] of each target computer, and *the target computer vulnerability database 714 is updated for each target computer in accordance with whether a response is received or is not received from each target computer for each TCP service port scanned and using the known vulnerability database to obtain the vulnerability information for the particular TCP service ports* that are determined to be open.” (*McClure* at 31:28-36) It does not appear to Applicant that updating a target computer vulnerability database or using a known vulnerability database, as stated in Passage C, discloses, teaches, or suggests receiving



anything from a network intrusion detection system, let alone receiving from such a network intrusion detection system one or more *data packets associated with an alarm* indicative of a potential attack on a target host, as recited in Claim 1.

As another example, *McClure* fails to disclose, teach, or suggest “identifying characteristics of the alarm from the data packets, including at least an attack type and an operating system fingerprint of the target host,” “comparing the attack type to the operating system type” and “indicating whether the target host is vulnerable to the attack based on the comparison,” as recited in Claim 1. As allegedly disclosing all of these limitations, the Examiner cites column 17, line 29 through column 18, line 50 of *McClure*.

The cited portion of *McClure* appears to disclose a technique for identifying the operating system of a target computer. (See *McClure* at 18:43-44) For example, the cited portion of *McClure* discloses sending messages to a target computer and saving responses from the target computer as fingerprints. (*Id.* at 17:29-64) The fingerprints are then compared to a known database of fingerprints associated with various operating systems and operating system versions. (*Id.* at 17:65-68) According to *McClure*, known fingerprints can be compiled through application of the above methodology to various target computers known to have a particular operating system before testing. (*Id.* at 17:67-18:3) The remainder of the cited portion appears to disclose various additional details related to the technique for identifying the operating system disclosed in *McClure*, including updating of the operating system fingerprint database, types of operating system fingerprints, and the types of messages that may be sent to the target computer to obtain responses from the target computer. (*Id.* at 18:20-50) However, the cited portion of *McClure* does not appear to disclose, teach, or suggest “identifying characteristics of the alarm from the data packets, including at least an attack type and an operating system fingerprint of the target host,” “comparing the attack type to the operating system type” and “indicating whether the target host is vulnerable to the attack based on the comparison,” as recited in Claim 1.

For at least these reasons, Appellant submits that *McClure* fails to disclose, teach, or suggest each and every limitation recited in Claim 1. Therefore, Appellant respectfully submits

that Claim 1 and its dependent claims are patentable over *McClure* and request that the Board overturn these rejections.

For at least certain analogous reasons, Appellant submits that *McClure* fails to disclose, teach, or suggest each and every limitation recited in independent Claims 7 and 16. Therefore, Appellant respectfully submits that Claims 7 and 16 and their dependent claims are patentable over *McClure* and request that the Board overturn the rejections of these claims.

**B. Dependent Claims 3, 9, and 18 are Allowable**

Claims 3, 9, and 18 depend from independent Claims 1, 7, and 16, respectively, which are shown above to be patentable over *McClure*. Thus, for at least the reasons discussed above with respect to Claims 1, 7, and 16, Appellant respectfully submits that Claims 3, 9, and 18 are patentable over *McClure* and requests that the Board overturn the rejections of these claims.

Additionally, Claims 3, 9, and 18 recite further patentable distinctions over *McClure*. Claim 3, which Appellant discusses as an example, recites the following:

The computerized method of Claim 1, further comprising:  
monitoring a dynamic configuration protocol server;  
detecting that a lease issue has occurred for a new target host;  
accessing a storage location;  
determining whether an operating system fingerprint for the new target host already exists in the storage location; and  
if the operating system fingerprint for the new target host does exist,  
then purging the existing operating system fingerprint for the new target host from the storage location.

Claims 9 and 18 recite certain substantially similar limitations for systems.

The Examiner cites column 22, lines 32-67 as allegedly disclosing all of the limitations recited in Claim 3. (Final Office Action at 4) The cited portion of *McClure* states the following:

A. Host Discovery

As described in more detail below, the host discovery phase applies one, two or three distinct tests to each IP address on the scan list. Preferably, the scan

list is scanned in batches, where each batch of IP addresses is scanned in parallel (as described in more detail below) to identify host computers (i.e., live target computers).

i. First Test (ICMP Ping Request)

In a first host discovery test, a standard ICMP ping request is sent to each target computer. If a response is received, the target computer is removed from the scan list and placed on the live list. In one embodiment, this entails sending out an ICMP echo request packet to each host. Multiple ICMP packets can advantageously be sent out in parallel to more than one IP address in a batch. Typically, the system waits until an ICMP echo reply is received from all IP addresses in the batch or the ICMP echo request is timed out. As a result of this process, for each batch of IP addresses on the scan list, a list of IP addresses that responded to the ICMP echo request is removed from the scan list and placed on the live list.

ii. Second Test (Sending TCP Packets)

If no response is received from one or more IP addresses on the list in the first test, a set of TCP packets (either single SYN packets or full TCP connection sequences ("TCP full connect")) are sent to the remaining target computers in the scan list as a second host discovery test. Specifically, a list of "TCP discovery ports" is selected in one embodiment. The selection is based on the TCP ports that are most likely to be open. The TCP discovery port list is advantageously relatively short, and preferably includes well known service ports such as HTTP (hypertext transfer protocol), SMTP (simple mail transfer protocol) and the like. One non-exclusive example embodiment of a TCP host discovery list is shown in Table 3.

At a minimum, it does not appear that the cited portion of *McClure* discloses, teaches or suggests the following limitations recited in Claim 3:

- detecting that a lease issue has occurred for a new target host;
- determining whether an operating system fingerprint for the new target host already exists in the storage location; and
- if the operating system fingerprint for the new target host does exist, then purging the existing operating system fingerprint for the new target host from the storage location.

For example, the cited portion of *McClure* does not appear to mention any "lease issue," let alone "determining that a lease issue has occurred for a new target host," as recited in Claim 3. It is not even clear what teaching from this cited portion the Examiner is attempting to equate to this limitation.

As another example, the cited portion of *McClure* does not appear to disclose, teach, or suggest "determining whether an operating system fingerprint for the new target host already

exists in the storage location,” as recited in Claim 3. Instead, it appears that the cited portion relates to discovering which target host computers on a “scan list” are “live” using an ICMP ping request technique or a TCP packet technique. However, the cited portion does not appear to make any reference to any “operating system fingerprint,” let alone “determining whether an operating system fingerprint for the new target host already exists in the storage location,” as recited in Claim 3.

As another example, the cited portion of *McClure* does not appear to disclose, teach, or suggest “if the operating system fingerprint for the new target host does exist, then purging the existing operating system fingerprint for the new target host from the storage location,” as recited in Claim 3. First, at least because the cited portion of *McClure* does not disclose, teach, or suggest “determining whether an operating system fingerprint for the new target host already exists in the storage location,” the cited portion of *McClure* necessarily fails to disclose, teach, or suggest “if the operating system fingerprint for the new target host does exist, then purging the existing operating system fingerprint for the new target host from the storage location,” as recited in Claim 3. Second, as discussed above, it appears that the cited portion relates to discovering which target host computers on a “scan list” are “live” using an ICMP ping request technique or a TCP packet technique. However, the cited portion does not appear to make any reference to any “operating system fingerprint,” let alone “if the operating system fingerprint for the new target host does exist, then purging the existing operating system fingerprint for the new target host from the storage location,” as recited in Claim 3.

For at least these reasons, Appellant respectfully submits that Claims 3, 9, and 18 are patentable over *McClure* and request that the Board overturn the rejections of these claims.

**C. Dependent Claims 4, 10, and 19 are Allowable**

Claims 4, 10, and 19 depend from independent Claims 1, 7, and 16, respectively, which are shown above to be patentable over *McClure*. Thus, for at least the reasons discussed above with respect to Claims 1, 7, and 16, Appellant respectfully submits that Claims 4, 10, and 19 are patentable over *McClure* and requests that the Board overturn the rejections of these claims.

Additionally, Claims 4, 10, and 19 recite further patentable distinctions over *McClure*.

Claim 4, which Appellant discusses as an example, recites the following:

The computerized method of Claim 1, further comprising:  
monitoring a dynamic configuration protocol server;  
detecting that a lease expire has occurred for an existing target host;  
accessing a storage location;  
determining whether an operating system fingerprint for the existing target host already exists in the storage location; and  
if the operating system fingerprint for the existing target host does not exist, then disregarding the lease expire; and  
if the operating system fingerprint for the existing target host does exist, then purging the existing operating system fingerprint for the existing target host from the storage location.

Claims 10 and 19 recite certain substantially similar limitations for systems.

The Examiner again cites column 22, lines 32-67 as allegedly disclosing all of the limitations recited in Claim 4. (Final Office Action at 4-5) Appellant respectfully directs the Board's attention to Section III.C above for the text of the cited portion of *McClure*.

At a minimum, it does not appear that the cited portion of *McClure* discloses, teaches or suggests the following limitations recited in Claim 4:

- detecting that a lease expire has occurred for an existing target host;
- determining whether an operating system fingerprint for the existing target host already exists in the storage location;
- if the operating system fingerprint for the existing target host does not exist, then disregarding the lease expire; and
- if the operating system fingerprint for the existing target host does exist, then purging the existing operating system fingerprint for the existing target host from the storage location.

For example, the cited portion of *McClure* does not appear to mention any "lease issue," let alone "detecting that a lease expire has occurred for an existing target host," as recited in Claim 4. It is not even clear what teaching from this cited portion the Examiner is attempting to equate to this limitation.

As another example, the cited portion of *McClure* does not appear to disclose, teach, or suggest "determining whether an operating system fingerprint for the existing target host

already exists in the storage location,” as recited in Claim 4. Instead, it appears that the cited portion relates to discovering which target host computers on a “scan list” are “live” using an ICMP ping request technique or a TCP packet technique. However, the cited portion does not appear to make any reference to any “operating system fingerprint,” let alone “determining whether an operating system fingerprint for the existing target host already exists in the storage location,” as recited in Claim 4.

As another example, the cited portion of *McClure* does not appear to disclose, teach, or suggest “if the operating system fingerprint for the existing target host does not exist, then disregarding the lease expire,” as recited in Claim 4. First, at least because the cited portion of *McClure* does not disclose, teach, or suggest “determining whether an operating system fingerprint for the existing target host already exists in the storage location,” the cited portion of *McClure* necessarily fails to disclose, teach, or suggest “if the operating system fingerprint for the existing target host does not exist, then disregarding the lease expire,” as recited in Claim 4. Second, as discussed above, it appears that the cited portion relates to discovering which target host computers on a “scan list” are “live” using an ICMP ping request technique or a TCP packet technique. However, the cited portion does not appear to make any reference to any “operating system fingerprint,” let alone “if the operating system fingerprint for the existing target host does not exist, then disregarding the lease expire,” as recited in Claim 4.

As another example, the cited portion of *McClure* does not appear to disclose, teach, or suggest “if the operating system fingerprint for the existing target host does exist, then purging the existing operating system fingerprint for the existing target host from the storage location,” as recited in Claim 4. First, at least because the cited portion of *McClure* does not disclose, teach, or suggest “determining whether an operating system fingerprint for the existing target host already exists in the storage location,” the cited portion of *McClure* necessarily fails to disclose, teach, or suggest “if the operating system fingerprint for the existing target host does exist, then purging the existing operating system fingerprint for the existing target host from the storage location,” as recited in Claim 4. Second, as discussed above, it appears that the cited portion relates to discovering which target host computers on a “scan list” are “live” using an ICMP ping request technique or a TCP packet technique. However, the cited portion does not appear to make any reference to any “operating system fingerprint,” let alone “if the operating

system fingerprint for the existing target host does exist, then purging the existing operating system fingerprint for the existing target host from the storage location,” as recited in Claim 4.

For at least these reasons, Appellant respectfully submits that Claims 4, 10, and 19 are patentable over *McClure* and request that the Board overturn the rejections of these claims.

**D. Dependent Claims 5 and 20 are Allowable**

Claims 5 and 20 depend from independent Claims 1 and 16, respectively, which are shown above to be patentable over *McClure*. Thus, for at least the reasons discussed above with respect to Claims 1 and 16, Appellant respectfully submits that Claims 5 and 20 are patentable over *McClure* and requests that the Board overturn the rejections of these claims.

Additionally, Claims 5 and 20 recite further patentable distinctions over *McClure*. Claim 5, which Appellant discusses as an example, recites the following:

The computerized method of Claim 1, further comprising:  
after receiving the data packets, determining whether a format for the alarm is valid; and  
if the format is not valid, then disregarding the alarm; otherwise  
if the format is valid, then continuing the computerized method with the identifying characteristics step.

Claim 20 recites certain substantially similar limitations for a system.

The Examiner cites column 23, lines 26-52 as allegedly disclosing all of the limitations recited in Claim 5. (Final Office Action at 5) The cited portion of *McClure* states the following:

In one embodiment, a standard TCP SYN packet is sent to some or all of the ports on the TCP host discovery list for each target IP address (target computer.) As with the prior ICMP ping test, multiple IP addresses are advantageously tested in parallel (i.e., in batches) in a preferred embodiment. If a target computer responds with a TCP SYN ACK, then the target computer is added to the live list. Otherwise, the TCP SYN request to the target times out (i.e., a maximum time period passes without a response from the target computer).

In an alternative embodiment of the TCP scan test, a standard TCP full connect request initiated using the standard Window.RTM. Winsock interface. If the operating system confirms that a TCP three-way handshake has been

completed, then the target computer is added to the live list. If the target responds with a TCP RST ACK, an ambiguous response, the target computer is added to the "potentially live" list. Otherwise, the TCP request to the target times out.

The foregoing tests result in a list of live target computers (IP addresses) on the live list. The target computers on the live list are removed from the scan list. If there are any IP addresses that have not been confirmed on the "live list" or "potentially live list," then a third step of scanning selected UDP ports on the target computer is performed for IP addresses remaining on the scan list.

It does not appear to Appellant that the cited portion discloses, teaches, or suggests the limitations recited in Claim 5. The cited portion does not even mention an alarm, let alone "after receiving the data packets, determining whether a format for the alarm is valid," as recited in Claim 5. Moreover, at least because the cited portion of *McClure* does not disclose, teach, or suggest the cited portion of *McClure* further does not appear to disclose, teach, or suggest "determining whether a format for the alarm is valid," the cited portion necessarily fails to disclose, teach, or suggest disregarding the alarm of the format is not valid or continuing the computerized method with the identifying characteristics step if the format is valid, as recited in Claim 5. In fact, the cited portion does not appear to disclose, teach, or suggest taking any actions with respect to any determination regarding an alarm. Instead, it appears that the cited portion relates to "pinging" target computers to discover which target computers are live.

For at least these reasons, Appellant respectfully submits that Claims 5 and 20 are patentable over *McClure* and request that the Board overturn the rejections of these claims.

**E. Dependent Claim 13 is Allowable**

Claim 13 depends from independent Claim 7, which is shown above to be patentable over *McClure*. Thus, for at least the reasons discussed above with respect to Claim 7 (by way of Claim 1), Appellant respectfully submits that Claim 13 is patentable over *McClure* and requests that the Board overturn the rejections of this claim.

Additionally, Claim 13 recites further patentable distinctions over *McClure*. Claim 13 recites the following:



The system of Claim 7, wherein the software program has no access to the protected network.

The Examiner cites column 24, lines 50-67 of *McClure* as allegedly disclosing all of the limitations recited in Claim 13. (Final Office Action at 6) The cited portion of *McClure* states the following:

Unlike the data in traditional UDP port detection packets, the data contained within the UDP packets sent in accordance with the present invention are specifically designed to prompt a reply from the scanned host (i.e., target computer) based on knowledge of a service typically associated with the UDP port. If no information is available about the UDP port, standard data (for example, the data representing a simple ASCII character return or control character) are placed in a UDP packet. In one embodiment, an exemplary UDP data probe packet is designed to solicit a response from a NetBIOS name service that typically runs on UDP protocol at port 137. An exemplary UDP data probe for UDP port 137 is shown in Table 5. In this case, the probe is advantageously a NetBIOS node status request, which typically results in a known response from the UDP port.

Nowhere does this cited portion disclose, teach, or suggest that “the software program has no access to the protected network,” as recited in Claim 13. First, such the cited portion does not appear to make any reference to a software program that has no access to a protected network. Second, the software program recited in Claim 13 is the software program originally-recited in independent Claim 7. According to Claim 7, that software program is operable to: receive the one or more data packets; identify characteristics of the alarm from the data packets, including at least an attack type and an operating system fingerprint of the target host; identify the operating system type from the operating system fingerprint; compare the attack type to the operating system type; and indicate whether the target host is vulnerable to the attack based on the comparison. The cited portion of *McClure* does not appear to disclose, teach, or suggest a software program that is operable to perform each of these recited limitations, nor does the cited portion of *McClure* appear to refer to such a software program, and that has no access to the protected network, as recited in Claim 13.

For at least these reasons, Appellant respectfully submits that Claim 13 is patentable over *McClure* and request that the Board overturn the rejections of this claim.


**Conclusion**

Appellant has demonstrated that, for at least the foregoing reasons, the present invention, as claimed, is clearly patentable over the references cited by the Examiner. Therefore, Appellant respectfully requests the Board to reverse the final rejection of the Examiner and instruct the Examiner to issue a Notice of Allowance of all pending claims.

The Commissioner is hereby authorized to charge the large entity fee of \$510.00 under 37 C.F.R. §§1.191(a) and 1.17(b) for filing this Appeal Brief to Deposit Account No. 02-0384 of Baker Botts L.L.P. Although no other fees are believed to be due at this time, the Commissioner is hereby authorized to charge any necessary additional fees and/or credit any overpayments to Deposit Account No. 02-0384 of Baker Botts L.L.P.

Respectfully submitted,

BAKER BOTTS L.L.P.  
Attorneys for Appellant

  
Chad D. Terrell  
Reg. No. 52,279

**Date:** December 5, 2007

Customer Number: **05073**

**Appendix A: The Claims**

1. (Rejected) A computerized method for reducing the false alarm rate of network intrusion detection systems, comprising:

receiving, from a network intrusion detection sensor, one or more data packets associated with an alarm indicative of a potential attack on a target host;

identifying characteristics of the alarm from the data packets, including at least an attack type and an operating system fingerprint of the target host;

identifying the operating system type from the operating system fingerprint;

comparing the attack type to the operating system type; and

indicating whether the target host is vulnerable to the attack based on the comparison.

2. (Rejected) The computerized method of Claim 1, further comprising storing the operating system fingerprint of the target host in a storage location for a time period.

3. (Rejected) The computerized method of Claim 1, further comprising:  
monitoring a dynamic configuration protocol server;  
detecting that a lease issue has occurred for a new target host;  
accessing a storage location;  
determining whether an operating system fingerprint for the new target host already exists in the storage location; and

if the operating system fingerprint for the new target host does exist, then purging the existing operating system fingerprint for the new target host from the storage location.

A.2

4. (Rejected) The computerized method of Claim 1, further comprising:  
monitoring a dynamic configuration protocol server;  
detecting that a lease expire has occurred for an existing target host;  
accessing a storage location;  
determining whether an operating system fingerprint for the existing target host already exists in the storage location; and  
if the operating system fingerprint for the existing target host does not exist, then disregarding the lease expire; and  
if the operating system fingerprint for the existing target host does exist, then purging the existing operating system fingerprint for the existing target host from the storage location.
5. (Rejected) The computerized method of Claim 1, further comprising:  
after receiving the data packets, determining whether a format for the alarm is valid;  
and  
if the format is not valid, then disregarding the alarm; otherwise  
if the format is valid, then continuing the computerized method with the identifying characteristics step.
6. (Rejected) The computerized method of Claim 1, further comprising  
automatically alerting a network administrator if the target host is vulnerable to the attack.

A.3

7. (Rejected) A system for reducing the false alarm rate of network intrusion detection systems, comprising:

a network intrusion detection system operable to transmit one or more data packets associated with an alarm indicative of a potential attack on a target host;

a software program embodied in a computer readable medium, the software program, when executed by a processor, operable to:

receive the one or more data packets;

identify characteristics of the alarm from the data packets, including at least an attack type and an operating system fingerprint of the target host;

identify the operating system type from the operating system fingerprint;

compare the attack type to the operating system type; and

indicate whether the target host is vulnerable to the attack based on the comparison.

8. (Rejected) The system of Claim 6, further comprising a storage location operable to store the operating system fingerprint of the target host for a time period.

9. (Rejected) The system of Claim 7, wherein the software program is further operable to:

monitor a dynamic configuration protocol server;

detect that a lease issue has occurred for a new target host;

access a storage location;

determine whether an operating system fingerprint for the new target host already exists in the storage location; and

if the operating system fingerprint for the new target host does exist, then the software program is further operable to purge the existing operating system fingerprint for the new target host from the storage location.

A.4

10. (Rejected) The system of Claim 7, wherein the software program is further operable to:

- monitor a dynamic configuration protocol server;
- detect that a lease expire has occurred for an existing target host;
- access a storage location;
- determine whether an operating system fingerprint for the existing target host already exists in the storage location; and
- if the operating system fingerprint for the existing target host does not exist, then disregard the lease expire; and
- if the operating system fingerprint for the existing target host does exist, then purge the existing operating system fingerprint for the existing target host from the storage location.

11. (Rejected) The system of Claim 7, wherein the software program is further operable to automatically alert a network administrator of the attack if the target host is vulnerable to the attack.

12. (Rejected) The system of Claim 7, wherein the software program has no knowledge of the protected network architecture.

13. (Rejected) The system of Claim 7, wherein the software program has no access to the protected network.

14. (Rejected) The system of Claim 7, wherein the NIDS is vendor independent.

15. (Rejected) The system of Claim 7, wherein the NIDS does not support passive operating system fingerprinting.

A.5

16. (Rejected) A system for reducing the false alarm rate of network intrusion detection systems, comprising:

means for receiving, from a network intrusion detection sensor, one or more data packets associated with an alarm indicative of a potential attack on a target host;

means for identifying characteristics of the alarm from the data packets, including at least an attack type and an operating system fingerprint of the target host;

means for identifying the operating system type from the operating system fingerprint;

means for comparing the attack type to the operating system type; and

means for indicating whether the target host is vulnerable to the attack based on the comparison.

17. (Rejected) The system of Claim 16, further comprising means for storing the operating system fingerprint of the target host for a time period.

18. (Rejected) The system of Claim 16, further comprising:

means for monitoring a dynamic configuration protocol server;

means for detecting that a lease issue has occurred for a new target host;

means for accessing a storage location;

means for determining whether an operating system fingerprint for the new target host already exists in the storage location; and

if the operating system fingerprint for the new target host does exist, then means for purging the existing operating system fingerprint for the new target host from the storage location.

A.6

19. (Rejected) The system of Claim 16, further comprising:  
means for monitoring a dynamic configuration protocol server;  
means for detecting that a lease expire has occurred for an existing target host;  
means for accessing a storage location;  
means for determining whether an operating system fingerprint for the existing target host already exists in the storage location; and  
if the operating system fingerprint for the existing target host does not exist, then means for disregarding the lease expire; and  
if the operating system fingerprint for the existing target host does exist, then means for purging the existing operating system fingerprint for the existing target host from the storage location.

20. (Rejected) The system of Claim 16, further comprising:  
after receiving the data packets, means for determining whether a format for the alarm is valid; and  
if the format is not valid, then means for disregarding the alarm.

21. (Rejected) The system of Claim 16, further comprising means for automatically alerting a network administrator if the target host is vulnerable to the attack.



ATTORNEY DOCKET NO.  
062891.1166

PATENT APPLICATION  
USSN 10/685,726

**Appendix B: *McClure***



US007152105B2

(12) **United States Patent**  
**McClure et al.**

(10) **Patent No.:** **US 7,152,105 B2**  
(45) **Date of Patent:** **Dec. 19, 2006**

(54) **SYSTEM AND METHOD FOR NETWORK  
VULNERABILITY DETECTION AND  
REPORTING**

(75) Inventors: **Stuart C. McClure**, Ladera Ranch, CA (US); **George Kurtz**, Coto de Caza, CA (US); **Robin Keir**, Mission Viejo, CA (US); **Marshall A. Beddoe**, San Clemente, CA (US); **Michael J. Morton**, Anaheim Hills, CA (US); **Christopher M. Prosise**, Mission Viejo, CA (US); **David M. Cole**, Huntington Beach, CA (US); **Christopher Abad**, San Francisco, CA (US)

(73) Assignee: **McAfee, Inc.**, Santa Clara, CA (US)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 733 days.

(21) Appl. No.: **10/050,675**

(22) Filed: **Jan. 15, 2002**

(65) **Prior Publication Data**

US 2003/0195861 A1 Oct. 16, 2003

(51) **Int. Cl.**

**G06F 15/173** (2006.01)

**G06F 11/00** (2006.01)

**G05B 19/00** (2006.01)

(52) **U.S. Cl.** ..... **709/224; 726/25; 340/5.53**

(58) **Field of Classification Search** ..... **709/224; 726/23, 25; 340/5.53**

See application file for complete search history.

(56) **References Cited**

**U.S. PATENT DOCUMENTS**

5,881,236 A \* 3/1999 Dickey ..... 709/221  
5,892,903 A 4/1999 Klaus ..... 395/187.01  
5,931,946 A 8/1999 Terada et al. .... 713/201  
6,266,774 B1 7/2001 Sampath et al. .... 713/201

6,282,546 B1 8/2001 Gleichauf et al. .... 707/102  
6,298,445 B1 10/2001 Shostack et al. .... 713/201  
6,301,668 B1 10/2001 Gleichauf et al. .... 713/201  
6,324,656 B1 11/2001 Gleichauf et al. .... 714/37  
6,725,046 B1 \* 4/2004 Park ..... 455/450  
7,000,247 B1 2/2006 Banzhof ..... 726/2  
2001/0034847 A1 10/2001 Gaul, Jr. .... 713/201  
2002/0100036 A1 \* 7/2002 Moshir et al. .... 717/173  
2003/0014664 A1 \* 1/2003 Hentunen ..... 713/200  
2003/0101353 A1 \* 5/2003 Tarquini et al. .... 713/200  
2004/0117478 A1 \* 6/2004 Triulzi et al. .... 709/224  
2004/0187032 A1 \* 9/2004 Gels et al. .... 713/201

**OTHER PUBLICATIONS**

Graig Smith and Christopher Abad, Know Your Enemy: Passive Fingerprint, Sep. 3, 2001, pp. 1, 2 and 3.\*  
Declaration of Dan Kuykendall in Opposition to OSC Re: Preliminary Injunction, *Foundstone, Inc. v. NT OBJECTIVES, Inc.*, Case No. 02CC15350 (Sup. Ct. Cal.), dated Oct. 23, 2002, pp. 3-4.

(Continued)

*Primary Examiner*—Ario Etienne

*Assistant Examiner*—El Hadji M. Sall


(74) *Attorney, Agent, or Firm*—Zilka-Kotab, PC; Christopher J. Hamaty

(57)

**ABSTRACT**

A system and method provide comprehensive and highly automated testing of vulnerabilities to intrusion on a target network, including identification of operating system, identification of target network topology and target computers, identification of open target ports, assessment of vulnerabilities on target ports, active assessment of vulnerabilities based on information acquired from target computers, quantitative assessment of target network security and vulnerability, and hierarchical graphical representation of the target network, target computers, and vulnerabilities in a test report. The system and method employ minimally obtrusive techniques to avoid interference with or damage to the target network during or after testing.

**20 Claims, 15 Drawing Sheets**



## OTHER PUBLICATIONS

Declaration of Micahel J. Morton in Opposition to OSC Re: Preliminary Injunction, *Foundstone, Inc. v. NT OBJECTives, Inc.*, Case No. 02CC15350 (Sup. Ct. Cal.), dated Oct. 23, 2002, pp. 2-5.

Declaration of Eric Caso in Opposition to OSC Re: Preliminary Injunction, *Foundstone, Inc. v. NT OBJECTives, Inc.*, Case No. 02CC15350 (Sup. Ct. Cal.), dated Oct. 23, 2002, pp. 2-5 & Exhibits B, C and D.

Declaration of Jassen D. Glaser in Opposition to OSC Re: Preliminary Injunction, *Foundstone, Inc. v. NT OBJECTives, Inc.*, Case No. 02CC15350 (Sup. Ct. Cal.), dated Oct. 23, 2002, pp. 6-10 & Exhibits F, G.

Memorandum of Points and Authorities in Support of Defendants' Opposition to Plaintiff's OSC Re: Preliminary Injunction, *Foundstone, Inc. v. NT OBJECTives, Inc.*, Case No. 02CC15350 (Sup. Ct. Cal.), dated Oct. 23, 2002, pp. 9-10.

Dan Kuykendall, a/k/a Seek3r, "Legal Docs: The response to their complaint", <[www.kuykendall.org](http://www.kuykendall.org)>, printed Oct. 25, 2002, pp. 2-4.

Declaration of Stuart McClure in Support of Plaintiff's Application for Temporary Restraining Order and Preliminary Injunction, *Foundstone, Inc. v. NT OBJECTives, Inc.*, Case No. 02CC15350 (Sup. Ct. Cal.), dated Oct. 4, 2002, pp. 2-5.

Plaintiff's Memorandum of Points and Authorities in Support of Application for Temporary Restraining Order, Preliminary Injunction, *Foundstone, Inc. v. NT OBJECTives, Inc.*, Case No. 02CC15350 (Sup. Ct. Cal.), dated Oct. 4, 2002, pp. 9-13.

Plaintiff's Reply in Support of Motion for Preliminary Injunction, *Foundstone, Inc. v. NT OBJECTives, Inc.*, Case No. 02CC15350 (Sup. Ct. Cal.), dated Oct. 25, 2002, pp. 5-8.

Plaintiff Foundstone's Supplemental Reply in Support of Order to Show Cause Re: Preliminary Injunction, *Foundstone, Inc. v. NT OBJECTives, Inc.*, Case No. 02CC15350 (Sup. Ct. Cal.), dated Oct. 28, 2002, pp. 2-3.

Marshall Beddoe and Christopher Abad, *The Siphon Project: An Implementation of Stealth Target Acquisition & Information Gathering Methodologies*, Blackhat USA Conference 2001, presented Jul. 11, 2001\*, available at <<http://www.blackhat.com/html/bh-usa-01/bh-usa-01-speakers.html>>.

Craig Smith and Peter Grundl, *Know Your Enemy: Passive Fingerprinting*, Sep. 3, 2001\*, available at <<http://www.project.honeynet.org/papers/finger/>>.

Fyodor, *Remote OS detection via TCP/IP Stack Fingerprinting (NMAP)*, Oct. 18, 1998\*, available at <<http://www.insecure.org/nmap/nmap-fingerprinting-article.html>> and <<http://project.honeynet.org/papers/finger/traces.txt>>.

Fyodor, *The Art of Port Scanning*, Phrack Magazine, vol. 7, Issue 51, Sep. 1, 1997\*, available at <<http://www.insecure.org/nmap/p51-11.txt>>.

The International Search Report from PCT/US02/01093 mailed Aug. 5, 2002.

The International Preliminary Examination Report from PCT/US02/01093 mailed Oct. 27, 2003.

\* cited by examiner

FIG. 1

FIG. 2





FIG. 3






FIG. 5







FIG. 6







FIG. 8



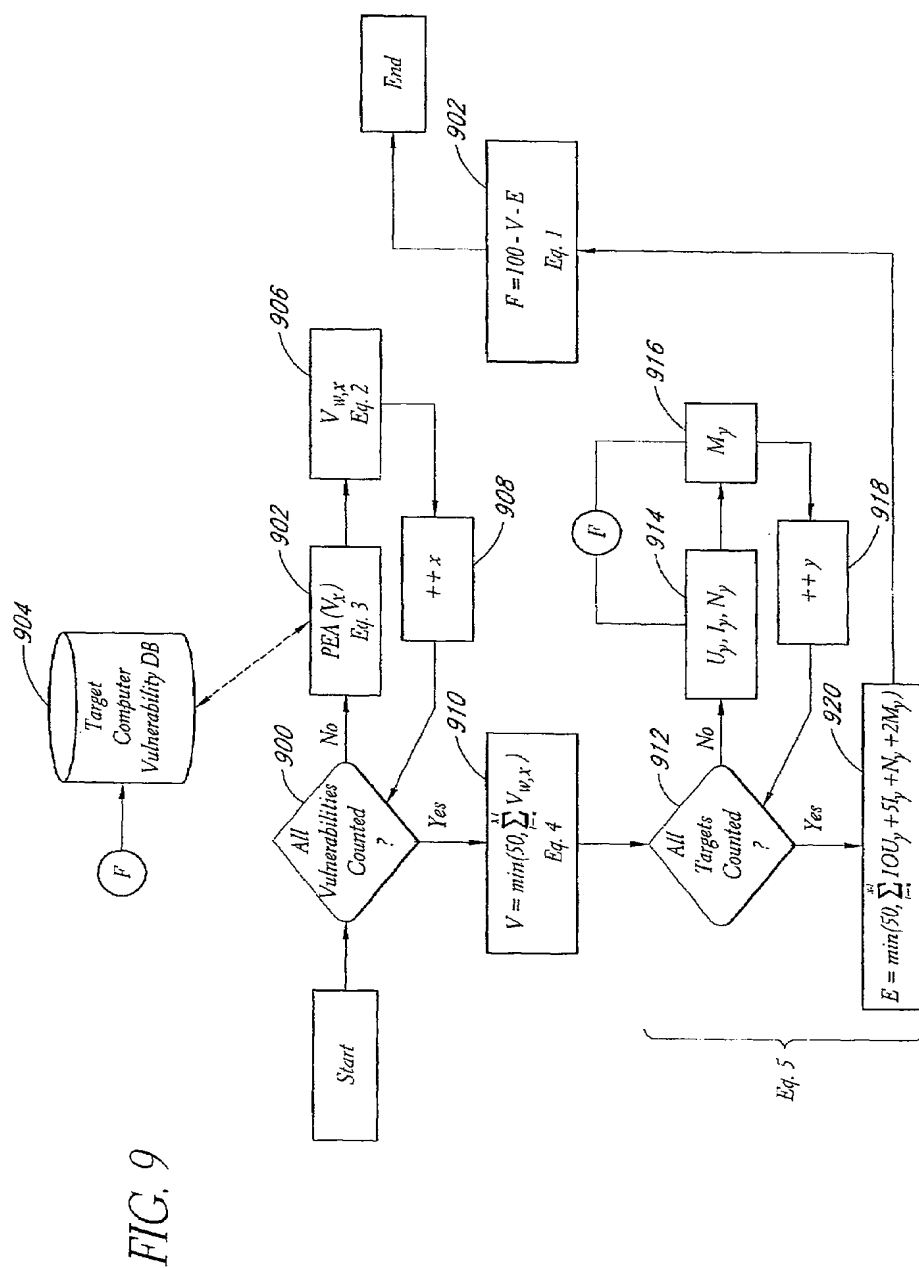

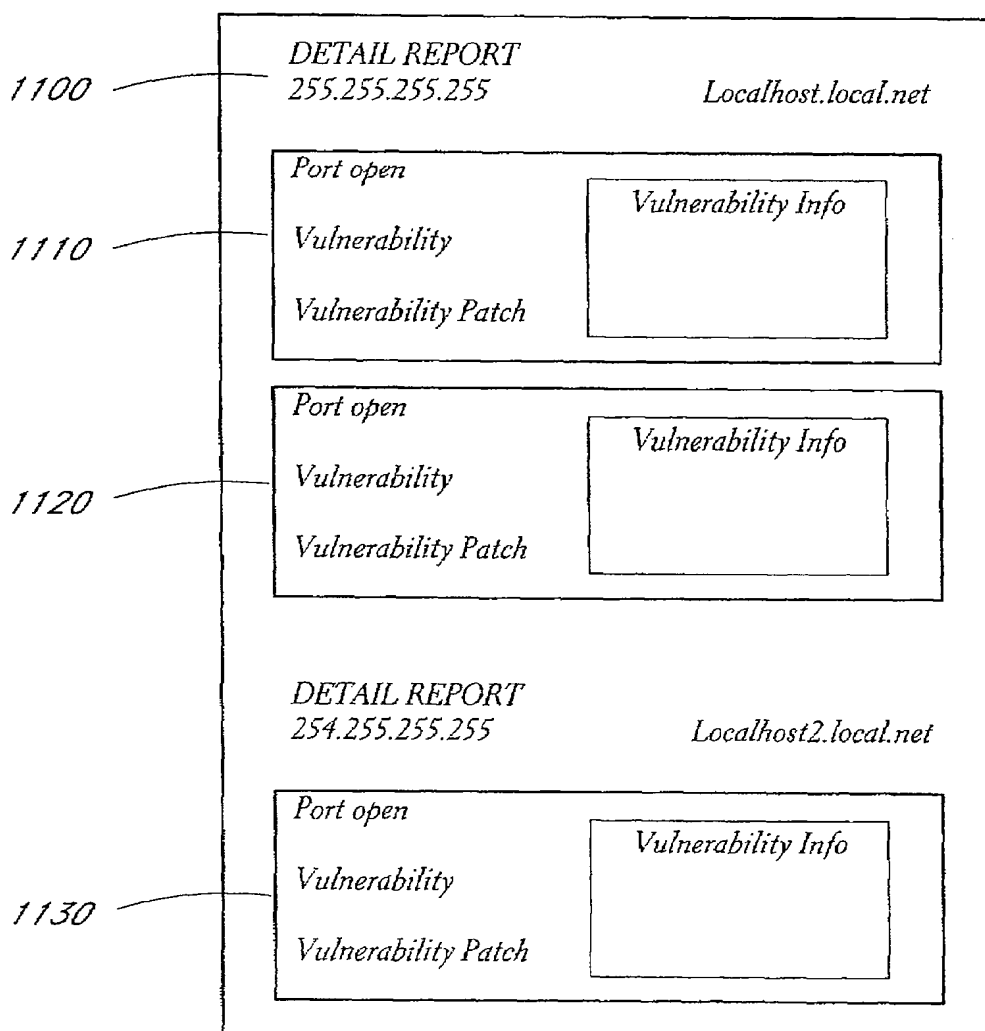
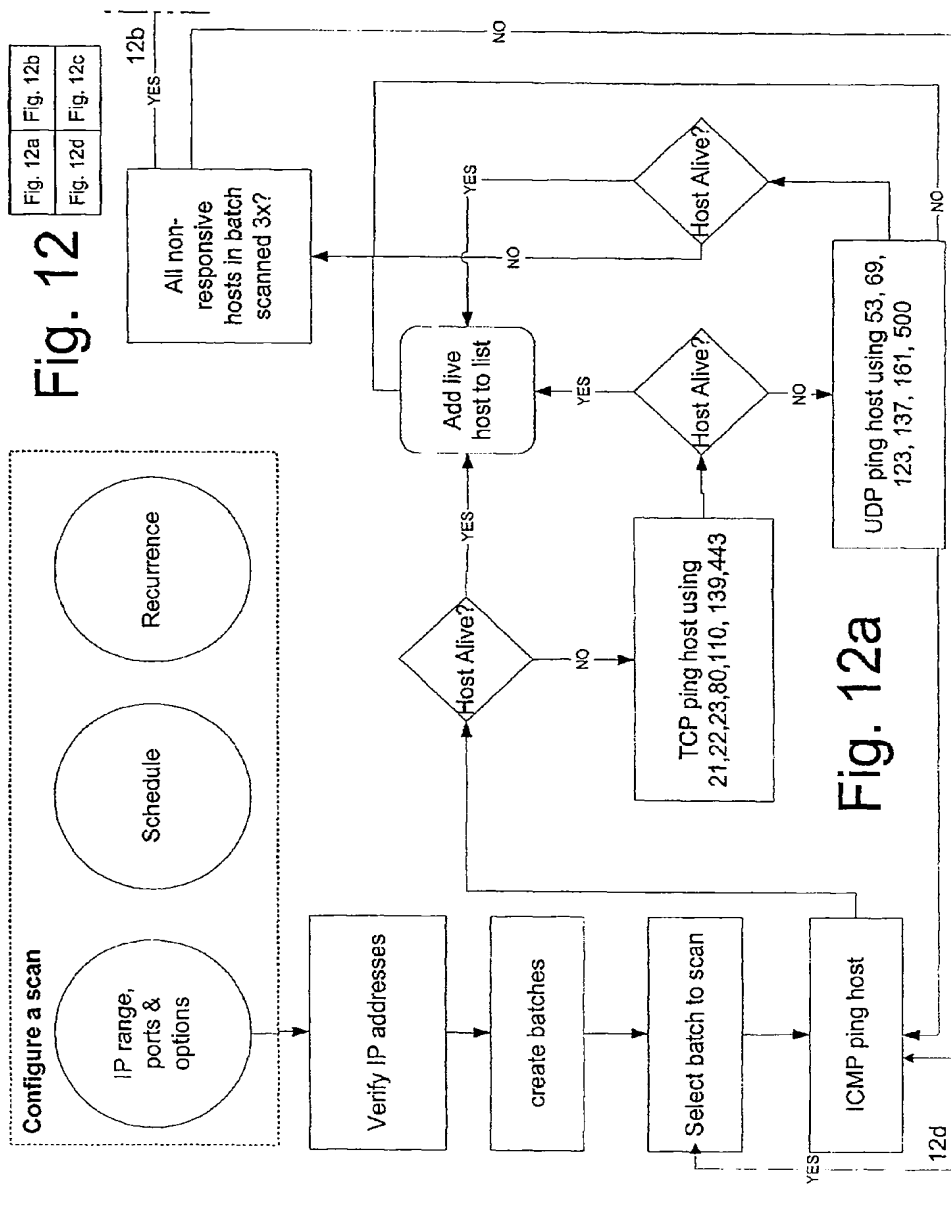
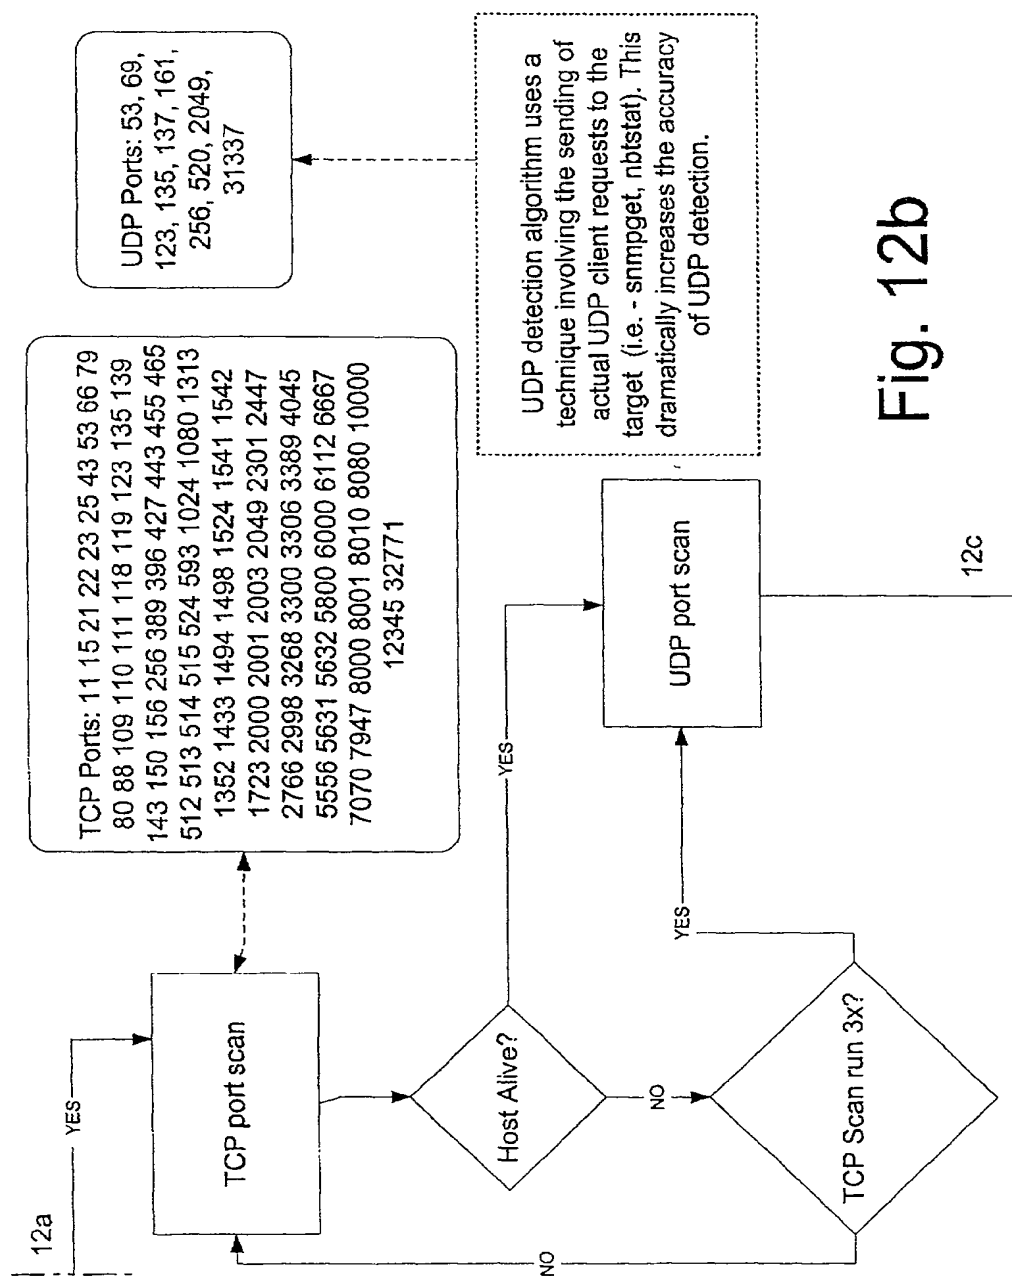


FIG. 10



*FIG. 11*







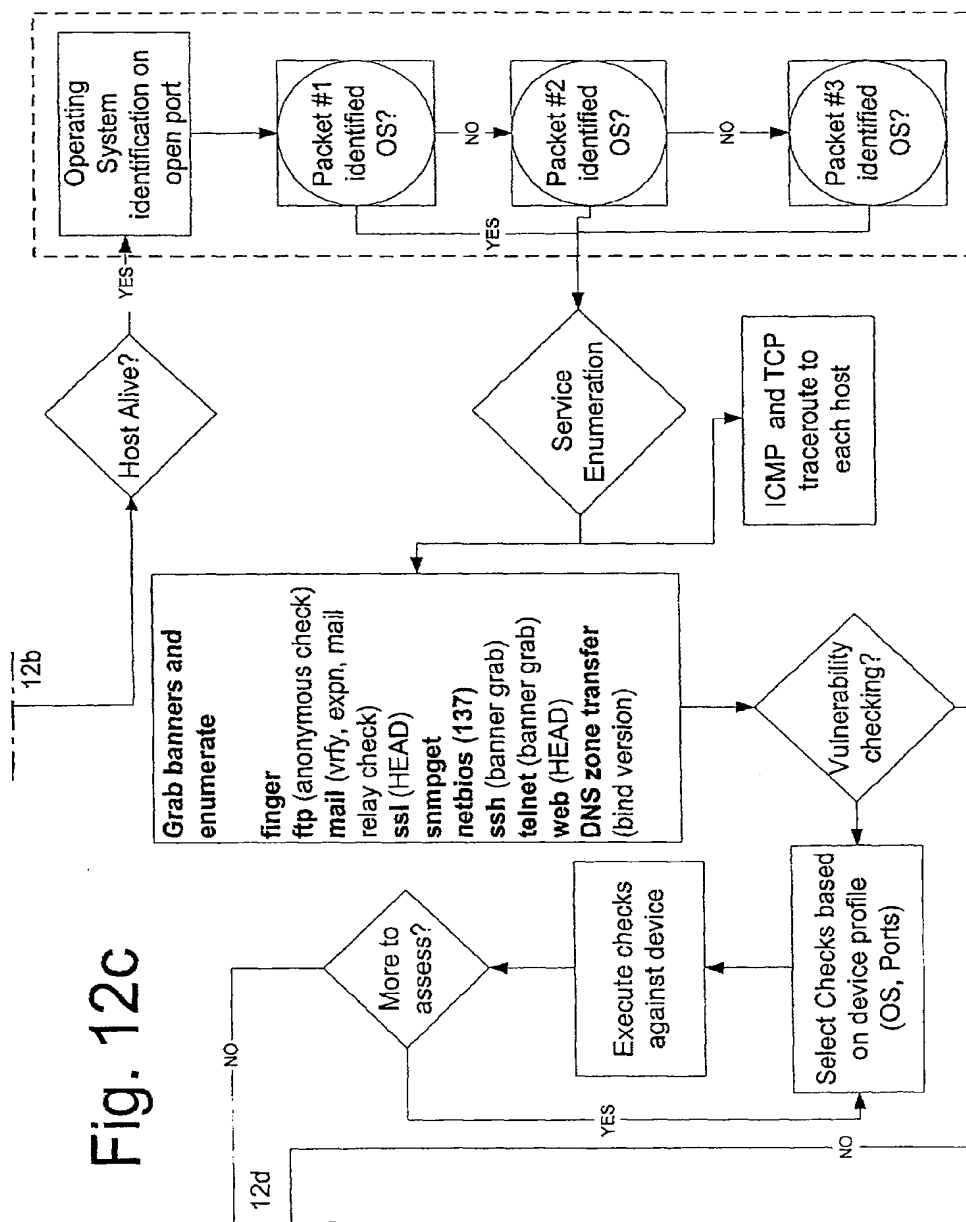



Fig. 12d



# SYSTEM AND METHOD FOR NETWORK VULNERABILITY DETECTION AND REPORTING

## BACKGROUND OF THE INVENTION

### 1. Field of the Invention

This invention relates to network system security, and more particularly relates to systems and methods for automatic detection, monitoring and reporting of network vulnerabilities.

### 2. Description of the Related Art

The reliability and security of a network is essential in a world where computer networks are a key element in intra-entity and inter-entity communications and transactions. Various tools have been used by network administrators, government, security consultants, and hackers to test the vulnerabilities of target networks, such as, for example, whether any computers on a network can be accessed and controlled remotely without authorization. Through this intensive testing, a target network can be "hardened" against common vulnerabilities and esoteric attacks. Existing testing systems, however, produce inconsistent results, use techniques that are unproven or that damage the target network, fail to respond to changing network environments or to detect new vulnerabilities, and report results in difficult to understand, text-based reports.

Well-known network security tools now exist to test network paths for possible intrusion. From a testing point, simple commands such as traceroute and ping can be used to manually map a network topography, and determine roughly what network addresses are "alive" with a computer "awake" on the network (i.e., determine which computers are on and are responding to network packets). A tool such as a port scanner can be used to test an individual target computer on the target network to determine what network ports are open. If open ports are found, these ports may provide access for possible intrusion, and potentially represent a vulnerability that can be exploited by a malicious hacker.

Some suites combining various network tools attempt to follow a quasi-automated process to test target computers on a target network. These suites provide variations on the tools described above, and provide long-form text-based output based on the outcome of this testing. The output of these security tests are extremely technical, and require extensive knowledge of network communications in order to interpret and provide advice based on the results. Thus, these partially automated suites do not provide comprehensive security to an entity seeking to "harden" its network.

Further, some security suites actually risk substantial damage to the target network. For example, while the use of malformed network packets to test a target computer can provide extensive information from the target and feedback on the security of the target, these malformed packets can destabilize the target computer in unpredictable ways. This sometimes results in a short-term loss of information to the target computer or, in more serious cases, a complete crash of the target computer operating system or hardware.

In other cases, the testing method used by existing suites is not reliable. If a network port scanning method employed on a target computer is, for example, 80% accurate over time, then a complete test of all 216 ports on a single computer may result in approximately 13,000 ports incorrectly identified as potentially running vulnerable services.

Over an entire target network, such "false positives" make it virtually impossible to determine the true security level of the target network.

Existing testing methods lack a standard, quantitative method for objectively comparing the security of a target network or target computer to other systems. Typically, a target network or target computer is ranked only as "high risk," "medium risk," or "low risk." However, such a three-tier system alone provides very little substantive feedback or comparative information about changes in the network over time, the relative weight of different vulnerabilities in determining the resulting risk level, or objective assessments of network security among otherwise heterogeneous network environment.

## SUMMARY OF THE INVENTION

The present invention solves these problems and more through a comprehensive network vulnerability testing and reporting method and system. Specifically, the testing system features include a selected combination of: (1) a non-destructive identification of target computer operating system; (2) a multiple-tier port scanning method for determination of what network addresses are active and what ports are active at those addresses; (3) a comparison of collected information about the target network with a database of known vulnerabilities; (4) a vulnerability assessment of some vulnerabilities on identified ports of identified target computers; (5) an active assessment of vulnerabilities reusing data discovered from previously discovered target computers; (6) an application of a quantitative score to objectively and comparatively rank the security of the target network; and, (7) reduction of detailed results of the information collected into hierarchical, dynamic and graphical representations of the target network, target computers, and vulnerabilities found therein. Other features are foreseen and disclosed herein, as well.

In its preferred embodiment, the testing system operates over a modem multi-layer packet network such as a corporate intranet or the Internet. The network typically includes one or more computers, where a computer includes a desktop station running any operating system, a router, a server, and/or any other networked device capable of sending and receiving packets through standard internet protocols such as TCP/IP (Transmission Control Protocol/Internet Protocol), UDP (User Datagram Protocol), and the like. The system and method can be run remotely from a monitoring computer outside the target network, or can be run by a monitoring computer included within the target network. The target network itself is typically defined as an interconnected set of computers, bounded by a specific pre-designated sub-network address, range of IP addresses or sub-addresses, physical network boundaries, computer names or unique identifiers, presence or connection via a pre-determined network protocol, and the like. The target computers comprise all or a portion of the computers found within the target network. For example, a target computer with a simple connection to a WAN (Wide Area Network) can be tested remotely, as a single peer target network. In a more complicated example, a distributed network provider can have multiple sub-networks geographically distributed throughout the world but interconnected via an internal protocol, as a WAN target network with thousands of target computers.

A target network typically runs on one or more IP-based network protocols. Most commonly, the protocol will be TCP/IP and UDP. Similarly, the testing system is typically indifferent to the physical layer structure and topology of the

target network. Only structural elements such as firewalls or routers that block, reroute, or change packets will affect the testing system. The testing system, however, attempts to adapt to these structural elements and generally provides accurate results regardless of physical implementation.

TCP/IP is a fundamental protocol used for packet-based network communications on local area networks, wide area networks, and global telecommunications networks such as the Internet. A sample configuration of a TCP/IP SYN (synchronization) packet is shown in Table 1.

TABLE 1

Typical TCP SYN packet					
Source Port		Destination Port			
Sequence Number					
Acknowledgement Number					
Data Offset	Reserved Data		Flags	Window	
Checksum		Urgent Pointer		Padding	
Options					
Data					

A computer typically runs on one or more operating systems. More commonly, these operating systems include those provided by Microsoft®, such as the Microsoft Windows® family of operating systems, MacOS® from Apple®, various flavors of UNIX including Linux®, NetBSD, FreeBSD, Solaris®, and the like. Additionally, devices on the target network may include router operating systems, mobile communication device operating systems, palmtop or handheld operating systems, appliance operating systems, set-top box operating systems, gaming operating systems, digital rights management systems, surveillance systems, smart card transaction systems, transportation management systems, and the like, that assign unique or temporary network addresses and are capable of sending and/or receiving traffic from the target network.

Target computers, in one embodiment, are identified by a unique or temporarily unique IP (Internet Protocol) address, typically in the form A.B.C.D, where each of A, B, C and D represent the Class A, Class B, Class C and Class D sub-networks and each has a value between 0 and 255. Typically, the target network is defined by one or more ranges of IP addresses controlled by the target network, but may contain additional target computers or target sub-networks connected to the target network topographically but not part of the predetermined IP range or ranges.

UDP (User Datagram Protocol) is an alternative "connectionless" communications protocol that runs above IP (Internet Protocol). UDP lacks the error correction and receipt acknowledgment features of connection-based protocols such as TCP. ICMP (Internet Control Message Protocol) is another extension of IP which permits control communications (most commonly through a ICMP PING request) between hosts on an IP network.

Another aspect of the invention includes non-destructive and relatively non-intrusive identification of the target operating system of a target computer.

Another aspect of the invention includes parallel testing of multiple target computers on a target network.

Another aspect of the invention includes an improved testing method to determine whether particular target computers on a target network are alive.

Another aspect of the invention includes an improved method for determining whether a set of commonly used ports are open on a target computer.

Another aspect of the invention includes an improved method for reliably determining whether a set of commonly used UDP ports are open or closed on a target computer.

Another aspect of the invention includes a method for associating the ports found open on a target computer with a known set of vulnerabilities.

Another aspect of the invention includes parallel testing of multiple ports and multiple target computers simultaneously.

Another aspect of the invention includes active assessment of some known set of vulnerabilities at a target computer.

Yet another aspect of the invention includes application of an objective quantitative score to the vulnerabilities found on a target network.

Still another aspect of the invention includes compilation of a dynamic, graphical report representing the network topology, network computers, and network vulnerabilities in a hierarchical report including both overview and detail documents.

In one embodiment, the present invention is a system for determining an operating system of a target computer operably connected to a network. The system comprises (1) first and second data packets, the first and second data packets compliant with a protocol supported by the network, the first and second data packets transmitted via the network to the target computer; (2) first and second operating system fingerprints comprising data bits stored in a computer-readable medium, the first and second operating system fingerprints associated with a first operating system; (3) a first target computer fingerprint comprising data bits stored in a computer-readable medium, the first target computer fingerprint including a representation of at least a portion of data received in response to the transmission of the first data packet; (4) a second target computer fingerprint comprising data bits stored in a computer-readable medium, the second target computer fingerprint including a representation of at least a portion of data received in response to the transmission of the second data packet; and (5) fingerprint comparison instructions executable by a computer to compare the first operating system fingerprint and the first target computer fingerprint, to compare the second operating system fingerprint and the second target computer fingerprint, and to generate a result indicative of whether the first operating system was running on the target computer. In a preferred aspect, the invention further comprises: (6) a third data packet, the third data packet compliant with the protocol, the first range of bits of the third data packet representing a third parameter value different from the first and second parameter values, the third data packet transmitted via the network to the target computer; (7) a third operating system fingerprint comprising data bits stored in a computer-readable medium, the third operating system fingerprint associated with the first operating system, the third operating system fingerprint differing from the first and second operating system fingerprints; and (8) a third target computer fingerprint comprising data bits stored in a computer-readable medium, the third target computer fingerprint including a representation of at least a portion of data received in response to the transmission of the first data packet, the comparison instructions executable by a computer to compare the third operating system fingerprint and the third target computer fingerprint before generating the result. In a further preferred aspect, the invention further comprises: (9) fourth, fifth and sixth operating system fingerprints comprising data bits stored in a computer-readable medium, the fourth, fifth and sixth operating system fingerprints associ-

5

ated with a second operating system, at least one of the fourth, fifth and sixth operating system fingerprints differing from a respective one of the first, second and third operating system fingerprints; the comparison instructions executable by a computer to compare the fourth operating system fingerprint and the first target computer fingerprint, to compare the fifth operating system fingerprint and the second target computer fingerprint, to compare the sixth operating system fingerprint and the third target computer fingerprint, and to generate a second result indicative of whether the second operating system was running on the target computer. Preferred aspects of this embodiment are ones wherein (10) the first parameter value is obtained by setting no bits, the second parameter value is obtained by setting one bit, and the third parameter value is obtained by setting two bits, or (11) wherein the first parameter value is 0, the second parameter value is 128, and the third parameter value is 128 plus a multiple of 256.

In another embodiment, the present invention is a system for determining an operating system of a target computer accessible via a network. The system comprises: (1) a plurality of data packets compliant with a protocol supported by the network, the plurality of data packets transmitted via the network to the target computer; (2) a first plurality of operating system fingerprints, each comprising data bits stored in a computer-readable medium, each associated with a first operating system; (3) a plurality of target computer fingerprints, each comprising data bits stored in a computer-readable medium, each including a representation of at least a portion of data received in response to the transmission of the plurality of data packets; and (4) fingerprint comparison instructions executable by a computer to compare the first plurality of the operating system fingerprint and the plurality of the target computer fingerprints, and to generate a result indicative of whether the first operating system was running on the target computer. A preferred aspect of the embodiment is one wherein the protocol is TCP/IP. Another preferred aspect of the embodiment further comprises (5) a second plurality of operating system fingerprints, each comprising data bits stored in a computer-readable medium, each associated with a second operating system, the fingerprint comparison instructions comparing the second plurality of the operating system fingerprints and the plurality of the target computer fingerprints to generate a second result indicative of whether the second operating system was running on the target computer.

A further embodiment of the present invention is a method for determining an operating system of a target computer accessible via a network. The method comprises the steps of (1) transmitting to the target computer a plurality of data packets compliant with a protocol supported by the network; (2) generating a plurality of target computer fingerprints, each including at least a portion of data received via the network in response to the transmission of the plurality of data packets; (3) comparing the plurality of target computer fingerprints to a first set of predetermined operating system fingerprints, each of the first set of predetermined operating system fingerprints associated with a first operating system; and (4) generating a result indicative of whether the first operating system was running on the target computer. In a preferred aspect the embodiment comprises the further steps of (5) comparing the plurality of target computer fingerprints to a second set of predetermined operating system fingerprints, each of the second set of a predetermined operating system fingerprints associated with a second operating system; and (6) generating a result indicative of whether the second operating system was

6

running on the target computer. One preferred aspect of that embodiment is one wherein the protocol is TCP/IP and wherein the value of the MSS option of two of the plurality of data packets is divisible by 128. Another preferred aspect of that embodiment is one wherein a first of the plurality of data packets has a maximum segment size option of 0, wherein a second of the plurality of data packets has a maximum segment size option of 128, and wherein a third of the plurality of data packets has a maximum segment size option of 384.

A still further embodiment of the invention is a method for identifying an operating system of a target computer via a network, the method comprising the steps of: (1) sending a first data packet to the target computer via the network, the first data packet complying with a protocol of the network and having a first pattern of bits in a first range of bits; (2) generating a first response value representing at least a portion of data received via the network in response to the sending of the first data packet; (3) sending a second data packet to the target computer via the network, the second data packet complying with the protocol and having a second pattern of bits in a first range of bits, the second pattern of bits different from the first pattern; (4) generating a second response value representing at least a portion of data received via the network in response to the sending of the second data packet; (5) sending a third data packet to the target computer via the network, the third data packet complying with the protocol and having a third pattern of bits in a first range of bits, the third pattern of bits different from the first or the second pattern; (6) generating a third response value representing at least a portion of data received via the network in response to the sending of the third data packet; (7) comparing the first response value to a first predetermined value associated with a first operating system; (8) comparing the second response value to a second predetermined value associated with the first operating system; (9) comparing the third response value to a third predetermined value associated with the first operating system; and (10) generating a value indicative of a relationship between the first operating system and the target computer. A preferred aspect of the embodiment comprises the further steps of: (11) comparing the first response value to a fourth predetermined value associated with a second operating system; (12) comparing the second response value to a fifth predetermined value associated with the second operating system; and (13) comparing the third response value to a sixth predetermined value associated with the second operating system. A preferred aspect of that embodiment is one wherein no bit is set in the first pattern of bits, wherein one bit is set in the second pattern of bits, and wherein two bits are set in the third pattern of bits. Another preferred aspect of that embodiment is one wherein the number of bytes in the second pattern of bits that have at least one bit set is greater than the number of bytes in the first pattern of bits that have at least one bit set, and wherein the number of bytes in the third pattern of bits that have at least one bit set is greater than the number of bytes in the second pattern of bits that have at least one bit set.

Yet another embodiment of the present invention is a system for determining whether a target computer is on a network, the system comprising: (1) a first set of port identifiers stored in a computer-readable medium, each of the first set of port identifiers representing a port used by computers to receive data packets compliant with a first protocol of the network, each of the first set of port identifiers representing a port associated with known network services; (2) a first set of data packets, each directed to a port

represented by at least one of the first set of port identifiers, each of the first set of data packets compliant with the first protocol and transmitted to the target computer via the network; (3) a first set of acknowledgement packets received via the network in response to the transmission of the first set of data packets, and (4) a list of host identifiers, each host identifier representing a computer on the network that transmits data in response to a packet sent to the respective computer, a host identifier representing the target computer added to the list of host identifiers if the first set of acknowledgement packets indicates a responsiveness of the target computer. An alternative preferred aspect of the embodiment further comprises: (5a) a second set of port identifiers stored in a computer-readable medium, each of the second set of port identifiers representing a port used by computers to receive data packets compliant with a second protocol of the network, each of the second set of port identifiers representing a port associated with known network services; (6a) a second set of data packets, each directed to a port represented by at least one of the second set of port identifiers, each of the second set of data packets compliant with the second protocol and transmitted to the target computer via the network, at least one of the second set of data packets including data associated with the known network services; (7a) a second set of acknowledgement packets received via the network in response to the transmission of the second set of data packets; and (8a) a host identifier representing the target computer added to the list of host identifiers if the second set of acknowledgement packets indicates a responsiveness of the target computer. A preferred aspect of that embodiment is one wherein the first protocol is TCP, wherein the second protocol is UDP, wherein the second set of acknowledgement packets is a nonzero set of UDP data response packets. Another alternative preferred aspect of the embodiment further comprises: (5b) a second set of port identifiers stored in a computer-readable medium, each of the second set of port identifiers representing a port used by computers to receive data packets compliant with a second protocol of the network, each of the second set of port identifiers representing a port associated with known network services; (6b) a second set of data packets, each directed to a port represented by at least one of the second set of port identifiers, each of the second set of data packets compliant with the second protocol and transmitted to the target computer via the network, at least one of the second set of data packets including data associated with the known network services; (7b) a second set of acknowledgement packets received via the network in response to the transmission of the second set of data packets; and (8b) a host identifier representing the target computer added to a second list of host identifiers if the second set of acknowledgement packets does not indicate an unresponsiveness of the target computer, each of the second list of host identifiers representing a computer not known to be unresponsive. A preferred aspect of that embodiment is one wherein the first protocol is TCP, wherein the second protocol is UDP, wherein the second set of acknowledgement packets is an empty set of ICMP error packets. A further preferred aspect of either alternative embodiment further comprises: (9) a third set of data packets, each directed to a port represented by at least one of the second set of port identifiers, each compliant with the second protocol, the third set of data packets transmitted to the target computer throughout a predetermined maximum latency period; (10) a first response received first in time in response to the transmission of the third set of data packets; (11) a second response received second in time in response to the trans-

mission of the third set of data packets, a time duration between the receipt of the first response and the receipt of the second response defining a target computer latency period. A further preferred aspect of the embodiment is one wherein each of the second set of data packets is transmitted continuously to the target computer for the duration of the target computer latency period.

A still further embodiment of the present invention is a system for testing the accessibility of a target computer via a network. The system comprises: (1) a set of port identifiers stored in a computer-readable medium, each of the set of port identifiers representing a UDP-compliant port, at least one of the port identifiers representing a port associated with known network services; (2) a set of UDP-compliant data packets, each associated with a port represented by at least one of the set of port identifiers, each of the UDP-compliant data packets transmitted continuously to the target computer for a duration approximately the same as the latency period of the target computer, at least one of the UDP-compliant data packets including data associated with the known network services; (3) a first list representing computers accessible via the network, the first list including the target computer if a nonzero set of UDP data response packets is received in response to the transmission of the data packets; and (4) a second list representing computers not known to be inaccessible via the network, the second list including the target computer if an empty set of ICMP error packets is received in response to the transmission of the data packets.

Another embodiment of the present invention is a method for determining whether a target computer is accessible via a network. The method comprises the steps of: (1) identifying TCP ports; (2) sending first data packets to the TCP ports of the target computer, each of the first data packets compliant with TCP; (3) receiving first acknowledgment packets in response to the sending of the first data packets; and (4) adding a representation of the target computer to a list representing accessible computers if the first acknowledgment packets are nonzero. A preferred aspect of the embodiment comprises the further steps of: (5) identifying UDP ports associated with network services; (6) sending second data packets to the UDP ports of the target computer, at least one of the second data packets sent continuously to the target computer throughout a latency period of the target computer; (7) receiving second acknowledgment packets in response to the sending of the second data packets; and (8) adding a representation of the target computer to a list representing accessible computers if the second acknowledgment packets are nonzero UDP data response packets. A further preferred aspect of the embodiment comprises the further step of: (9) determining the latency period of the target computer by measuring the time between responses received in response to packets transmitted to the target computer. A further preferred aspect of the embodiment comprises the further step of: (10) adding a representation of the target computer to a list representing computers not known to be inaccessible via the network, the adding performed if the second acknowledgment packets comprise an empty set of ICMP error packets.

An additional embodiment of the present invention is a method for assessing the vulnerability of a target computer via a network. The method comprising the steps of: (1) discovering a set of responsive computers on a network by transmitting a set of ICMP packets, a set of TCP packets and a set of UDP packets to a group of computers on a network; (2) detecting services on each of the set of responsive computers by transmitting TCP packets to first ports of each of the set of responsive computers and by transmitting UDP

packets to second ports of each of the set of responsive computers, the first and second ports commonly used by computers to receive data packets over a network, the TCP packets including data associated with at least one computer-based service known to use one of the first ports, the UDP packets including data associated with at least one computer-based service known to use one of the second ports; and (3) generating a list of responsive ports using responses received in response to the transmission of the TCP packets and the UDP packets. A preferred aspect of the embodiment comprises the further step of: (4) determining an operating system used by each of the set of responsive computers by comparing predetermined values with portions of responses received from each of the set of responsive computers in response to transmission of a plurality of TCP-compliant packets to each of the set of responsive computers. A further preferred aspect of the embodiment comprises the further step of: (5) confirming the presence of vulnerabilities on the network by applying an automated vulnerability script to each responsive port represented in the list of responsive ports, each of the automated vulnerability scripts testing a vulnerability known to be associated with a computer configuration comprising a particular responsive port and a particular operating system. A still further preferred aspect of the embodiment comprises the further step of: (6) calculating an objective indicia of security of the network, the calculation based on a weighted summation of confirmed vulnerabilities. A preferred aspect of the embodiment comprises the further step of: (7) determining a topology of the network, the topology determination made by transmitting a set of ICMP packets with varying time to live (TTL) settings and by transmitting a set of TCP packets with varying TTL settings. Another preferred aspect of the embodiment comprises the further step of: (8) producing a graphical representation of the network, the representation including a topological map of the network, a color-based representation of weighted confirmed vulnerabilities, and an association between the graphical representation and information descriptive of confirmed vulnerabilities and computers on the network.

Another embodiment of the present invention is a method for creating a topological representation of a network. The method comprises the steps of: (1) identifying responsive computers on the network; (2) obtaining a plurality of sequences of IP addresses by sending to each responsive computer a sequence of packets having increasing TTL values, each sequence of IP addresses representing nodes in the network between a source computer and one of the responsive computers, adjacent IP addresses in each sequence representing connected nodes, each of the nodes comprising a computer or a router; (3) generating a list of node structures, each of the node structures including data representing a node and data indicative of other nodes to which it directly connects, the list representing all IP addresses in the plurality of sequences; (4) determining for each IP address a distance count, the distance count representing a number of nodes between a node having the IP address and a source node; (5) creating a router structure for each node structure that represents a node comprising a router; (6) associating with each of the router structures connection data representative of each connecting node that connects to no other node except the router represented by the respective router structure; (7) for each router structure, visually depicting a graphical shape spatially related to one or more graphical shapes corresponding to connecting nodes represented by the connection data of the respective router structure; and (8) for each router structure, visually depicting

a connection between a graphical shape associated with the respective router structure and another graphical shape associated with a different router structure when distance counts associated with the IP addresses of routers represented by the respective router structure and the different router structure indicate a direct connection. A preferred aspect of the embodiment comprises the further step of: (9) testing whether a router represented by a router structure and a connecting node represented in connection data comprise two network connections of one node. A further preferred aspect of this embodiment is one wherein the graphical shape representing a router is a sphere, and wherein each of the spatially related graphical shapes is a sphere orbiting the sphere representing the router.

Yet another embodiment of the present invention is a method for calculating an objective security score for a network. The method comprising the steps of: (1) determining a vulnerability value numerically representing a combination of known vulnerabilities of a network; (2) determining an exposure value numerically representing a combination of accessible ports of computers on the network; and (3) deriving a score by combining the vulnerability value and the exposure value. A preferred aspect of this embodiment is one wherein the combination of known vulnerabilities is a summation of weighted numeric expressions of particular vulnerabilities, the weighting based on an ease of exploitation ranking and on an access granted ranking for each vulnerability.

Still another embodiment of the present invention is a method for conducting an automated network vulnerability attack, the method comprising the steps of: (1) selecting a set of vulnerability attacks for each responsive computer on a network, each selected vulnerability attack for each responsive computer designed to expose a vulnerability associated with ports of the respective computer known to be accessible and also associated with an operating system used by the respective computer; (2) encoding the set of vulnerability attacks such that each is represented in a database by a unique identifier; (3) representing each of the set of vulnerability attacks using instructions of an automated scripting language; and (4) executing the vulnerability attacks by processing the instructions with a computer.

One more embodiment of the present invention is a hierarchical network vulnerability report. The report comprises: (1) a first report level comprising: (a) an objective score representing the security of the network; and (b) a graphical representation of a network topology, including a graphical representation of computers accessible via the network and a color-based graphical representation of the vulnerability of at least some of the computers; and (2) a second report level comprising: (a) a textual list describing the computers and their associated vulnerabilities; and (b) an exposure report describing accessible ports and services of the computers.

An additional embodiment of the present invention is a vulnerability assessment language. The vulnerability assessment language comprises: (1) a set of programming language statements used to create executable scripts, the scripts executed in a thread-safe execution architecture wherein all variables are stack variables and wherein a parse tree is treated as a read-only data structure; (2) a set of special scalar data types interchangeable with an integer data type in expressions, each of the set of special scalar data types having a set of constant values configured to support vulnerability assessment operations embodied in scripts; (3) a set of native objects declared in a metascope owning a script scope to make available the native objects to execut-

able scripts, the native objects facilitating network communication, providing callable member functions for building lists of unique ports and directing script execution to certain hosts, and providing IP addresses for scripts; and (4) a vulnerability object behaving to copy itself into a global data area where other scripts may access its information to compromise another machine, facilitating the use by one script of vulnerability data discovered by a different script.

A further embodiment of the present invention is a method for automated application of a known vulnerability on a target computer. The method comprises the steps of: (1) providing a database of known vulnerabilities, the database including a data object; (2) providing an executable script, the executable script associated with the data object; (3) applying the executable script to the target computer, the script performing the known vulnerability on a port of the target computer; and (4) returning a value representing at least one of the success, failure or other outcome of the executable script.

A still further embodiment of the present invention is a method for automated application of known vulnerabilities to target computers of a network. The method comprises the steps of: (1) providing a database of known vulnerabilities; (2) providing a set of executable scripts, each executable to apply a known vulnerability to a specified target computer; (3) executing first executable scripts to apply vulnerabilities on specified target computers; (4) monitoring return values representing a success, failure or other outcome of each of the first executable scripts; and (5) generating a report using the return values, the report representing a security level of the network. One preferred aspect of this embodiment comprises the further step of: (6) identifying execution time intervals wherein execution of the first executable scripts commences at the beginning of each of the time intervals and pauses at the end of each of the time intervals, until all of the first executable scripts have executed. A preferred aspect of the embodiment comprises the further step of: (7) automatically repeating the execution of the first executable scripts when the execution of the first executable scripts is completed. Another preferred aspect of the embodiment comprises the further steps of: (8) generating a report upon each completed execution of the first executable scripts; and (9) calculating a security trend for the network by comparing a plurality of the reports. An alternative preferred aspect of the embodiment comprises the further step of: (10) executing second executable scripts to apply vulnerabilities to a second network of computers during the execution of the first executable scripts. Another preferred aspect of the embodiment is one wherein the second network is a subset of the network. Still another preferred aspect of the embodiment is one wherein the first executable scripts are configured to apply vulnerabilities to a first port of all of the target computers before applying vulnerabilities to a second port of all of the target computers. An additional preferred aspect of the embodiment comprises the further step of allocating a plurality of packet slots, each packet slot permitting asynchronous transmission of a packet by one of the executable scripts.

#### BRIEF DESCRIPTION OF THE DRAWINGS

Preferred embodiments of the present invention are described below in connection with the attached drawings in which:

FIG. 1 illustrates one embodiment of a target network;

FIG. 2 illustrates one embodiment of a target computer on the target network;

FIG. 3 illustrates one embodiment of a comprehensive testing method;

FIG. 4 illustrates one embodiment of the operating system identification method;

FIG. 5 illustrates one example embodiment of the TCP SYN packet used in the operating system identification method of FIG. 3;

FIG. 6 illustrates one embodiment of first phase scanning to determine what target computers are alive;

FIG. 7 illustrates one embodiment of second phase scanning to determine what ports are open on a target computer;

FIG. 8 illustrates one embodiment of active assessment of a vulnerability of a target computer on a target network;

FIG. 9 illustrates one embodiment of a methodology for determining the security score for a target network;

FIG. 10 illustrates one embodiment of a hierarchical security report, including a graphical representation of network topology and network vulnerabilities; and

FIG. 11 illustrates a second embodiment of a hierarchical security report in greater detail.

FIG. 12 illustrates a second embodiment of the comprehensive testing method.

#### DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

##### I. Basic Implementation, Structure and Control Language

FIG. 1 illustrates one embodiment of a target network. The network security system 100 of the present invention is, in one embodiment, at least one Intel-based server running on a Windows 2000 operating system, although any computer system or operating system capable of handling an IP network and capable of large-scale data processing can be used. The network security system 100 may be outside the target network 102 or inside the target network (not shown). In either case, the system 100 is connected to the target network 102 through a network such as the Internet, via one or more nodes 104. The target network 102, in one example, consists of an intranet with a central intranet hub 106. The target network 102 further includes a firewall 108 which blocks some incoming traffic to or outgoing network traffic leaving the target network 102. The target network further comprises a number of hosts 110, defined as within a predetermined range of Internet Protocol (IP) addresses. In some cases, external hosts 112 may lie outside the target network but may nonetheless be connected to the target network 102.

FIG. 2 illustrates one embodiment of a target computer on the target network. In general, a host IP address represents a target computer, as more generally defined below, if the address is in use by the target network. In a simplified representation of a target computer 200 at a host 110, the target computer 200 is running an operating system 202. The operating system preferably contains at least one network TCP/IP stack 204 to provide packet transport, preferably including an interface to provide raw socket 206 connections between the target computer 200 and the network. The physical connection to the network 208 is provided, in one embodiment, by a Network Interface Card (NIC) 210. On an IP network, a packet can be received at any one of 65,536 logical ports 212 at the target computer 200. Similarly, any number of network services 214 can be provided.

FIG. 3 illustrates one embodiment of a comprehensive testing method in accordance with embodiments of the present invention. FIG. 3 presents an overview of the



method. Additional details of the method are set forth below in connection with FIGS. 4–11.

In a first step or routine 310 in FIG. 3, the testing method defines the target network and creates a scan list of IP addresses. The scan list is stored in a scan list database 312. Then, in a routine 314, the method obtains a first batch of IP addresses from the scan list database 312 and begins a detailed analysis of the target network by performing a host discovery routine that begins in a block 320.

The host discovery routine comprises an ICMP host discovery routine 322, which will be described in more detail below. During the ICMP host discovery routine 322, the method pings the target computers identified by the present batch of IP addresses. Based on the responses or lack of responses, the ICMP host discovery routine 322 is able to determine that certain target computers are “live” or “probably live,” and the associated IP addresses are added to a respective live database 324 or probably live database 326 accordingly.

After completing the ICMP host discovery routine 322, the method performs a decision routine 328 wherein the method determines whether all the IP addresses in the current batch of IP addresses have been discovered (i.e., whether all the IP addresses have been assigned to the live database 324 or the probably live database 326. If any IP address has not been discovered, the method proceeds to a TCP host discovery routine 330, which will be described in more detail below. During the TCP host discovery routine 330, the method sends TCP packets to the remaining target computers identified by the present batch of IP addresses. Based on responses or lack of responses, the TCP host discovery routine 330 is able to determine that certain ones of the remaining target computers are “live” or “probably live,” and the associated IP addresses are added to a respective live database 324 or probably live database 326 accordingly.

After completing the TCP host discovery routine 330, the method performs a decision routine 332 wherein the method determines whether all the IP addresses in the current batch of IP addresses have been discovered (i.e., whether all the IP addresses have been assigned to the live database 324 or the probably live database 326. If any IP address has not been discovered, the method proceeds to an intelligent UDP host discovery routine 334, which will be described in more detail below. During the intelligent UDP host discovery routine 334, the method sends UDP packets to the remaining target computers identified by the present batch of IP addresses. Based on responses or lack of responses, the intelligent UDP host discovery routine 334 is able to determine that certain ones of the remaining target computers are “live” or “probably live,” and the associated IP addresses are added to a respective live database 324 or probably live database 326 accordingly.

After completing the intelligent UDP host discovery routine 334, the method performs a decision routine 336 wherein the method determines whether all the IP addresses in the current batch of IP addresses have been discovered (i.e., whether all the IP addresses have been assigned to the live database 324 or the probably live database 326. If any IP address has not been discovered, the method, in one embodiment, proceeds to an intensive UDP host discovery routine 338, which will be described in more detail below. During the intensive UDP host discovery routine 338, the method sends additional UDP packets to the remaining target computers identified by the present batch of IP addresses. Based on responses or lack of responses, the intensive UDP discovery routine 338 is able to determine

that certain ones of the remaining target computers are “live” or “probably live,” and the associated IP addresses are added to a respective live database 324 or probably live database 326 accordingly. The intensive UDP host discovery routine 338 is optional and may not be included in all embodiments.

After completing the intensive UDP host discovery routine 338, the method preferably proceeds to a service discovery routine that begins in a block 340. Alternatively, in certain embodiments, the foregoing host discovery routines 322, 330, 334, 338 are advantageously repeated to determine whether additional IP addresses corresponding to target computers can be assigned to the live database 324 and the probably live database 326. In such alternative embodiments, the method repeats the host discovery routines a maximum of a predetermined times before continuing to the service discovery routine 340. Those IP addresses for which no response is received by any method are, in one embodiment, added to a dead list 339 of hosts.

In the service discovery routine, the method performs a TCP service discovery routine 342, which will be described in more detail below. In the TCP service discovery routine 342, the method sends TCP packets to selected ports of the discovered target computers in the live list 324 and the probably live list 326 and monitors the responses. Based on the responses or lack of responses, the TCP discovery routine 342 adds information regarding open ports of the target computers to a target computer database 344.

After performing the TCP service discovery routine 342, the method performs a UDP service discovery routine 346, which will be described in more detail below. In the UDP service discovery routine 346, the method sends UDP packets to selected ports of the discovered target computers in the live list 324 and the probably live list 326 and monitors the responses. Based on the responses or lack of responses, the UDP discovery routine 346 adds information regarding open ports of the target computers to the target computer database 344. In alternative embodiments, the TCP service discovery routine 342 and the UDP service discovery routine 346 are advantageously repeated a limited number of times to determine whether additional open ports are discovered.

After completing the UDP service discovery routine 346 (or after completing the routine a limited number of times), the method proceeds to an operating system (OS) identification routine 350 wherein the method determines the type and version of operating system present on each of the live computers having open ports. As will be described in more detail below, the method sends two or more (preferably three) RFC compliant TCP packets to the target computers. The TCP packets have predetermined data in at least one selected field. Each target computer responds to each of the TCP packets. The response information from each computer is treated as a “fingerprint” for each computer. The fingerprint from each computer is compared with fingerprints in fingerprint database 352, and the results of the comparisons are used to identify the target computers with high degrees of accuracy.

After completing the operating system identification routine 350 for the target computers, the method proceeds to a traceroute routine that begins in a block 354. In the traceroute routine, the method first performs an ICMP traceroute routine 356 in which the method uses ICMP tracerouting techniques, described in more detail below. In particular, the method sends a plurality of ICMP echo request packets to the target computer with varying TTL (time to live) values in the TCP/IP header. The method creates a network topology based on the known TTL value, the number of “hops” between the system and the target computer, and the router/

host at each "hop." The information from the ICMP traceroute routine 356 is added to a network map database 358. When all packets sent have either arrived, failed to arrive, or timed out, the traceroute step is completed for that target computer. The ICMP traceroute steps are repeated until a complete trace is received or a selected predetermined maximum number of passes is completed.

After completing the predetermined number of ICMP traceroute passes, the method proceeds to a decision routine 360 wherein the method determines whether the trace of the target computer is complete or incomplete. If the trace is incomplete, the method proceeds to a TCP traceroute routine 362. Otherwise, the method bypasses the TCP traceroute routine and proceeds directly to a vulnerability assessment routine 364.

As described in more detail below, the TCP traceroute routine 362 works similarly to ICMP traceroute routine 354, except that TCP SYN packets are sent to the target computers. As with ICMP tracerouting, the TTL value in each SYN packet is incrementally increased, and the return of ICMP unreachable packets and SYN ACK packets is monitored for all "hops" between the scanning system and the target host. Through the combination of ICMP tracerouting and TCP tracerouting, a complete map to each target computer, and collectively a relatively complete map of the target network topology is advantageously created and stored in the network map database 358.

After completing the TCP traceroute routine 362, the method proceeds to the vulnerability assessment routine 364. As described in more detail below, in the vulnerability assessment routine 364, the method executes vulnerability scripts that apply known vulnerabilities to open ports of the live target computers to determine whether the ports of the target computers exhibit the potential vulnerabilities. The method uses information stored in a known vulnerability database 366 to select the vulnerability scripts to execute for each active port. Information collected from vulnerable target computers is advantageously stored to the target computer database 344.

In one embodiment, the vulnerability assessment routine 364 preferably only performs vulnerability checks associated with the identified operating system and open ports of the target computer as determined by the operating system identification routine 350 and service discovery routine 340. If the operating system is not conclusively identified, typically the routine runs all known vulnerabilities for the open ports of the target computer.

In one embodiment, the method proceeds to an active assessment routine 365. As further described in detail below, the active assessment routine 365 uses information collected from target computers in the vulnerability assessment routine 364 to execute further vulnerability scripts to open ports of the live target computers. Specifically, the active assessment routine 365 reuses known vulnerabilities and information collected from target computers to continue to determine whether the ports of the target computers exhibit potential vulnerabilities, using information from the known vulnerability database 366 and information collected in the target computer database 344.

After completing the active assessment routine 365, the method proceeds to a decision routine 368 to determine whether the method has analyzed all the potential target computers. In particular, in the decision routine 368, the method determines whether the last batch of IP address has been analyzed. If further batches of target computers remain to be analyzed, the method proceeds to the routine 314 to

obtain the next batch of IP addresses. Otherwise, the method proceeds to a scoring routine 370.

In the scoring routine 370, described in more detail below, the method establishes a vulnerability score for each target computer and for the network based on the results of the active assessment and based on the vulnerability information in the known vulnerability database 366. The method then proceeds to a reporting routine 372, also described in more detail below, wherein the method reports the results of the scanning, active assessment and scoring.

The method can be advantageously repeated continuously on a target network, and can be scheduled into pre-determined scanning window time periods for execution of the method over time. For example, in preferred embodiments, the method is scheduled to scan during off-peak hours when the network is less likely to be less heavily used. In particularly preferred embodiments, the method is interruptible at the end of a window of off-peak hours and resumes where it paused at the beginning of the next off-peak window. In particularly advantageous embodiments, the method operates on multiple target networks concurrently by threading to share network resources.

## II. Non-Destructive Operating System Identification

Vulnerability of and access to a target computer may be heightened by knowing which particular operating system is running on the computer. Identifying a target computer's operating system can be accomplished by examining the operating system's response to a data packet it receives over the network. The forms of the packets and the responses to the packets are generated in accordance with network protocols. The written definitions of the protocols used for communications on the Internet are set forth in Internet Request for Comment (RFC) documents. For example, the TCP/IP protocol is defined in part in RFC 793, incorporated herein by reference, and contains a standard model for TCP packet communications over a network. While virtually every operating system includes a TCP/IP stack, each TCP/IP stack is implemented in a slightly different manner. Thus, known responses from a TCP/IP stack unique to a particular operating system can serve as a "fingerprint" to determine the target computer's actual operating system.

Operating system detection traditionally has been performed by sending out a combination of RFC-compliant and non-RFC compliant TCP packets. The traditional system then collects the unique, non-standard response from the target computer and maps the response to a known database of TCP/IP stacks and related operating systems. However, this method tends to be inaccurate, highly dependent on the particular packet shape and TCP/IP stack of the target computer, and requires a large number of packets to identify the target operating system with any degree of reliability. This method may trigger security or firewall alarms at the target computer, and, more seriously, this method may interfere with or actually crash the target computer, in large part due to the non-RFC compliant packets sent to the target.

The present system typically employs a unique set of new features to maximize the accuracy of operating system detection while minimizing intrusiveness and interference with operations of the target computer. In one embodiment, the invention sends RFC-compliant TCP "SYN" (synchronization) packets to a target computer. The use of RFC-compliant TCP packets advantageously reduces the probability that the detection packets are blocked by a router or firewall, and greatly reduces the probability that the detection packets will cause damage or crashes at the target computer. In one particularly preferred embodiment, the

invention uses just three RFC-compliant TCP packets. Thus, network strain is significantly reduced during detection of the operating systems of a large number of target computers on a target network.

In one embodiment, the first packet sent is a completely standard TCP SYN to an open port on the target computer.

compiled through application of the above methodology to various target computers known to have a particular operating system before testing. For example, testing of known computers running various versions of the Solaris® operating system provide the following operating system fingerprints:

TABLE 2

Sample OS Fingerprints for Solaris and BSD operating systems									
AW <sub>0</sub>	AW <sub>128</sub>	AW <sub>384</sub>	OPT <sub>0</sub>	OPT <sub>128</sub>	OPT <sub>384</sub>	TTL	DF	Flags	OS
83E8	8380	8400	02040218	02040080	02040180	FF	0	SA--	Solaris 2.6
6050	6000	6000	02040564	02040564	02040564	40	0	SA--	Solaris 2.7
6050	6000	6000	020405B4	020405B4	020405B4	40	0	SA--	Solaris 8
4000	4000	4080	---	---	---	64	0	---	OpenBSD 2.x
4000	4000	4000	---	---	---	64	0	---	NetBSD 1.4.2

The MSS (maximum segment size) option in the options field of the first packet is set to 0 (i.e., no bits set in the MSS option). See FIG. 5. When an acknowledgement packet, a SYN ACK packet, is received by the detection system from the target computer, certain bits from the packet are saved by the system. In one embodiment, for example, the TCP advertised window, TTL (time-to-live), options, flags, and the DF (don't fragment) fields are saved to a first fingerprint.

In this embodiment, a second packet is then sent. The second TCP SYN packet is again a standard TCP SYN packet; however, the MSS option is set to 128 in the second packet (corresponding to the setting of a single bit in the MSS option). Portions of the response SYN ACK from the target computer (preferably the TCP advertised window, TTL, and DF bits) are again saved to a second fingerprint. Finally, a third TCP SYN packet is sent. The third packet is also a standard TCP SYN packet; however, the MSS option is set to 384 in the third packet (corresponding to the setting of two bits in the MSS option). Again, portions of the response SYN ACK from the target computer (preferably the TCP advertised window, TTL, and DF bits) are once again saved to a third fingerprint.

In one embodiment, the fingerprint is saved in the following format:

AW<sub>MSS=0</sub>:AW<sub>MSS=128</sub>:AW<sub>MSS=384</sub>:TTL:DF:OS

where, for example,

AW=TCP advertised window

MSS=TCP Options Maximum Segment Size

TTL=TCP Options Time to Live

DF=TCP Options Don't Fragment Flag, and

OS=Operating System identification.

In another embodiment, the fingerprint is saved in the following format:

AW<sub>MSS=0</sub>:AW<sub>MSS=128</sub>:AW<sub>MSS=384</sub>:OPT<sub>MSS=0</sub>:OPT<sub>MSS=128</sub>:OPT<sub>MSS=384</sub>:TTL:DF:FL:OS

where, for example,

OPT=TCP Options Bytes, and

FL=TCP Flags.

The fingerprints are then compared to a known database of fingerprints associated with various operating systems and operating system versions. Known fingerprints can be

While more than one OS fingerprint may be associated with each operating system, collisions between fingerprints of distinct operating systems have been found to be highly unlikely. Tables can be compiled for other operating systems similar to that shown in Table 2. As operating system versions and popularity change over time, the fingerprint database is advantageously regularly updated to account for patches, version changes, and new operating systems. The fingerprint style shown above is only one embodiment of such a database, and any efficient method to store the operating system fingerprint can be used, based upon the TCP options altered, number of packets typically sent to a target computer, other TCP fields stored for recognition, and identification field used to represent a particular operating system and version, and the like. In one example, a unique data string for a particular operating system is compressed and stored using a digest algorithm such as MD5, and the like. For further example, perfect matching of fingerprints is not required: a system may employ a percentage match, such as, for example, 90% similarity between two fingerprints, as sufficient to identify a target computer as having a particular operating system or at least being in particular family of operating systems.

Below is an example exchange of packets when performing an OS identification. Three standard TCP SYN packets are sent to the remote host. The first packet is a SYN packet with no data and no IP or TCP options. Packet 2 is also a TCP SYN packet but a TCP Maximum Segment Size of 128 in the TCP options field is set. The third and final packet is again a TCP SYN packet but a TCP Maximum Segment Size of 384 is set in the TCP options field.

As noted above, by analyzing the replies from the 3 packets, a fingerprint is assembled that appears in a textual format as follows:

window1:window2:window3:options1:options2:options3:  
ttl:dontfrag:flags

where,

window1=the TCP window size received from the 1st response

window2=the TCP window size received from the 2nd response

window3=the TCP window size received from the 3rd response

options1=the option bytes received from the 1st response

options2=the option bytes received from the 2nd response

options3=the option bytes received from the 3rd response

ttl=the IP TTL field from the 1st response

don't frag= is the IP don't Fragment bit from the 1st response, and

flags= are single character representations of the TCP flags from 1st response.

In the example TCP packets shown below, the resultant fingerprint looks like this:

40E8:4000:4080:020405B4:020405B4:020405B4:80:1:SA

From this, the fingerprint is compared to a database of known operating system fingerprints in order to find the closest match that will identify the remote operating system. In this example, three example TCP packets sent and three example TCP packets returned are shown below:

---

Packet 1 send  
TCP  
Source port: 272  
Destination port: 80  
Sequence: 0x01100000 (17825792)  
Acknowledgement: 0x00000000 (0)  
Header length: 0x05 (5) - 20 bytes  
Flags: SYN  
URG: 0  
ACK: 0  
PSH: 0  
RST: 0  
SYN: 1  
FIN: 0  
Window: 0x0040 (64)  
Checksum: 0x4518 (17688) - correct  
Urgent Pointer: 0x0000 (0)  
TCP Options: None  
Data length: 0x0 (0)

Packet 1 reply  
TCP  
Source port: 80  
Destination port: 272  
Sequence: 0x659A2C81 (1704602753)  
Acknowledgement: 0x01100001 (17825793)  
Header length: 0x06 (6) - 24 bytes  
Flags: SYN ACK  
URG: 0  
ACK: 1  
PSH: 0  
RST: 0  
SYN: 1  
FIN: 0  
Window: 0xE840 (59456)  
Checksum: 0x9A47 (39495) - correct  
Urgent Pointer: 0x0000 (0)  
TCP Options  
Maximum Segment Size: 0x5B4 (1460)  
Data length: 0x0 (0)

Packet 2 send  
TCP  
Source port: 528  
Destination port: 80  
Sequence: 0x03100000 (51380224)  
Acknowledgement: 0x00000000 (0)  
Header length: 0x07 (7) - 28 bytes  
Flags: SYN  
URG: 0  
ACK: 0  
PSH: 0  
RST: 0  
SYN: 1  
FIN: 0  
Window: 0x0040 (64)  
Checksum: 0x1E8A (7818) - correct  
Urgent Pointer: 0x0000 (0)  
TCP Options  
Maximum Segment Size: 0x80 (128)  
Data length: 0x0 (0)

Packet 2 reply  
TCP  
Source port: 80

-continued

---

Destination port: 528  
Sequence: 0x659ABB23 (1704639267)  
Acknowledgement: 0x03100001 (51380225)  
Header length: 0x06 (6) - 24 bytes  
Flags: SYN ACK  
URG: 0  
ACK: 1  
PSH: 0  
RST: 0  
SYN: 1  
FIN: 0  
Window: 0x0040 (64)  
Checksum: 0x098D (2445) - correct  
Urgent Pointer: 0x0000 (0)  
TCP Options  
Maximum Segment Size: 0x5B4 (1460)  
Data length: 0x0 (0)

Packet 3 send  
TCP  
Source port: 784  
Destination port: 80  
Sequence: 0x05100000 (84934656)  
Acknowledgement: 0x00000000 (0)  
Header length: 0x07 (7) - 28 bytes  
Flags: SYN  
URG: 0  
ACK: 0  
PSH: 0  
RST: 0  
SYN: 1  
FIN: 0  
Window: 0x0040 (64)  
Checksum: 0x1A8A (6794) - correct  
Urgent Pointer: 0x0000 (0)  
TCP Options  
Maximum Segment Size: 0x180 (384)  
Data length: 0x0 (0)

Packet 3 reply  
TCP  
Source port: 80  
Destination port: 784  
Sequence: 0x659B732B (1704686379)  
Acknowledgement: 0x05100001 (84934657)  
Header length: 0x06 (6) - 24 bytes  
Flags: SYN ACK  
URG: 0  
ACK: 1  
PSH: 0  
RST: 0  
SYN: 1  
FIN: 0  
Window: 0x8040 (32832)  
Checksum: 0x4E04 (19972) - correct  
Urgent Pointer: 0x0000 (0)  
TCP Options  
Maximum Segment Size: 0x5B4 (1460)  
Data length: 0x0 (0)

---

While any number of iterations of the fingerprints described above can be derived for a target computer, it has been determined that three fingerprints provide the most accurate identification of operating system without undue repetitiveness. Similarly, while other TCP options flags may be altered to detect target operating systems, it has been found that alteration of the TCP advertised window, over a plurality of test SYN packets, is most effective, preferably with three test SYN packets with TCP option MSS values of 0, 128 and 384 respectively. Furthermore, alteration of the MSS value in the TCP advertised window is also less likely to interfere with target computer operation than alteration of other values in the packets.

The foregoing operating system identification method is summarized in FIG. 4. As described above, in accordance with the method of FIG. 4, a network security system first sends a first RFC compliant TCP SYN packet to a target

computer 412 via a first packet transmission represented by a line 420. The first TCP SYN packet has the TCP Options Maximum Segment Size (MSS) set to a value of 0 (i.e., all bits cleared). The target computer 412 responds to the first TCP SYN packet with a first SYN ACK packet represented by a line 422. As discussed above, at least a portion of the information included in the first SYN ACK packet received from the target computer 412 is determined by data in the TCP/IP stack within the target computer 412, and the data is determined, at least in part, by the particular operating system running on the target computer 412.

The network security system 410 next sends a second RFC compliant TCP SYN packet to a target computer 412 via a second packet transmission represented by a line 430. The first TCP SYN packet has the TCP Options Maximum Segment Size (MSS) set to a value of 128 (i.e., bit 7 set). The target computer 412 responds to the second TCP SYN packet with a second SYN ACK packet represented by a line 432. As discussed above, at least a portion of the information included in the second SYN ACK packet received from the target computer 412 is also determined by data in the TCP/IP stack within the target computer 412, and the data is determined, at least in part, by the particular operating system running on the target computer 412.

Preferably, the network security system 410 next sends a third RFC compliant TCP SYN packet to a target computer 412 via a third packet transmission represented by a line 440. The first TCP SYN packet has the TCP Options Maximum Segment Size (MSS) set to a value of 384 (i.e., bits 7 and 8 set). The target computer 412 responds to the third TCP SYN packet with a third SYN ACK packet represented by a line 442. As discussed above, at least a portion of the information included in the third SYN ACK packet received from the target computer 412 is also determined by data in the TCP/IP stack within the target computer 412, and the data is determined, at least in part, by the particular operating system running on the target computer 412.

Together, the information in the three SYN ACK packets received by the network security system 410 from the target computer 412 in response to the TCP SYN packets comprise a fingerprint that is compared with the fingerprint database 352 described above in connection with FIG. 3.

FIG. 5 illustrates one example embodiment of the TCP SYN packet 500 used in the operating system identification method of FIG. 3. On a higher (IP) level, the packet typically provides source and destination Internet Protocol addresses and unique network addresses (not shown). On the TCP level, the packet includes the source port 510 from which the packet was sent, and the destination port 512 on the target computer to which the packet is destined. The 32 bit sequence number 514 describes the starting point of the data contained in the packet in the data window, and the 32 bit acknowledgement number 516 contains the last byte received by the target computer. Data offset 518 and a reserved section 520 are also part of the packet.

The TCP flags 522 denote important information about the packet. In particular, SYN denotes the first packet in a new stream, and space in the sequences 514 and 516 is typically reserved for SYN flags besides the single bit in the TCP flags region 522. The window 524 describes how much data may be transferred into the buffer at one of the end point computers in the TCP packet communication. The checksum 526 and urgent pointer 528 are included. The TCP options 530 typically include a maximum segment size. After packet padding 532, the actual data 534 carried by the packet is attached.

### III. Multi-Tier Port Scanning for Target Computer Identification and Target Port Identification

Port scanning is an essential tool for ensuring network security. Typically, a would-be intruder will apply a port scanner to a chosen target computer to attempt to find open ports. Through these open ports, the intruder may hope to obtain access to the target computer through known or unknown vulnerabilities. In the network security context, applying an ordinary port scanner to all  $2^{16}$  (65,536) ports on each target computer on a target network may significantly drain network resources, take an impracticable amount of time, and not provide an accurate accounting of actual vulnerabilities of target computers.

In one embodiment, the present system employs an iterated port scanning system with at least two phases of operation: host discovery and service discovery. For a particular set of IP address ranges in the target network (the "scan list"), host discovery determines which IP addresses represent live target computers (i.e., computers that respond to network packets) and adds each such address to the "live list;" determines which IP addresses represent computers that are partially responsive, as discussed below, and adds each such address to the "potentially live list;" and determines which IP addresses represent computers that are non-responsive, and adds each such address to the "dead list." In the service scan, each target computer reported as live in the host discovery phase is subjected to scanning of a known set of ports likely to be open to traffic.

#### A. Host Discovery

As described in more detail below, the host discovery phase applies one, two or three distinct tests to each IP address on the scan list. Preferably, the scan list is scanned in batches, where each batch of IP addresses is scanned in parallel (as described in more detail below) to identify host computers (i.e., live target computers).

##### i. First Test (ICMP Ping Request)

In a first host discovery test, a standard ICMP ping request is sent to each target computer. If a response is received, the target computer is removed from the scan list and placed on the live list. In one embodiment, this entails sending out an ICMP echo request packet to each host. Multiple ICMP packets can advantageously be sent out in parallel to more than one IP address in a batch. Typically, the system waits until an ICMP echo reply is received from all IP addresses in the batch or the ICMP echo request is timed out. As a result of this process, for each batch of IP addresses on the scan list, a list of IP addresses that responded to the ICMP echo request is removed from the scan list and placed on the live list.

##### ii. Second Test (Sending TCP Packets)

If no response is received from one or more IP addresses on the list in the first test, a set of TCP packets (either single SYN packets or full TCP connection sequences ("TCP full connect")) are sent to the remaining target computers in the scan list as a second host discovery test. Specifically, a list of "TCP discovery ports" is selected in one embodiment. The selection is based on the TCP ports that are most likely to be open. The TCP discovery port list is advantageously relatively short, and preferably includes well known service ports such as HTTP (hypertext transfer protocol), SMTP (simple mail transfer protocol) and the like. One non-exclusive example embodiment of a TCP host discovery list is shown in Table 3.

TABLE 3

Sample TCP Discovery List	
Service	Port
FTP	21
SSH	22
Telnet	23
SMTP	25
HTTP	80, 81, 8000
POP3	110
NetBIOS	139
SSL/HTTPS	443
PPTP	1723

Other common ports can be added or removed from the list, and the list, for example, can be tailored to specific network environments where certain ports are more likely to be in use than others. In this example, File Transfer Protocol (FTP), Secure Shell (SSH), Telnet, Simple Mail Transfer Protocol (SMTP), HyperText Transfer Protocol (HTTP), Post Office Protocol (POP3), NetBios, Secure Sockets Layer (SSL), and Point-to-Point Tunneling Protocol (PPTP) are advantageously used.

In one embodiment, a standard TCP SYN packet is sent to some or all of the ports on the TCP host discovery list for each target IP address (target computer.) As with the prior ICMP ping test, multiple IP addresses are advantageously tested in parallel (i.e., in batches) in a preferred embodiment. If a target computer responds with a TCP SYN ACK, then the target computer is added to the live list. Otherwise, the TCP SYN request to the target times out (i.e., a maximum time period passes without a response from the target computer).

In an alternative embodiment of the TCP scan test, a standard TCP fill connect request initiated using the standard Window® Winsock interface. If the operating system confirms that a TCP three-way handshake has been completed, then the target computer is added to the live list. If the target responds with a TCP RST ACK, an ambiguous response, the target computer is added to the "potentially live" list. Otherwise, the TCP request to the target times out.

The foregoing tests result in a list of live target computers (IP addresses) on the live list. The target computers on the live list are removed from the scan list. If there are any IP addresses that have not been confirmed on the "live list" or "potentially live list," then a third step of scanning selected UDP ports on the target computer is performed for IP addresses remaining on the scan list.

### iii. Third Test (Intelligent UDP Port Scanning)

If any IP addresses (i.e., target computers) remain on the scan list after the first two tests, then a third test of intelligent UDP port scanning is performed on the remaining addresses. As described below, intelligent UDP port scanning differs from traditional scanning of UDP ports, which is notoriously inaccurate. When scanning TCP ports, a response from the target computer signals that the target port on the target computer is open. In contrast, while scanning UDP ports, no response from the target computer signals that the target port is open, and a response (an ICMP error message) will only be sent if the UDP port scanned is closed. The traditional method of scanning UDP ports thus results in a significant

number of "false positives" where a UDP scan results in no response (suggesting an open port), but the UDP port scanned is actually closed. This may happen, for example, when a firewall or router blocks the ICMP error message from returning from the target computer or when the ICMP error message is simply lost while returning from the target computer. Over thousands of tests, such errors become likely.

Sometimes, in order to "force" a response from the target computer, an intruder may send a malformed packet to a target port. While this known technique increases the likelihood that an open UDP port on the target computer can be identified, this technique also substantially increases the likelihood that the malformed packet could damage the target computer. Also, firewalls or routers may detect and filter out malformed packets, and such packets can alert the target network of an attempted security breach.

The intelligent UDP port scanning test in accordance with this embodiment of the present invention employs an efficient, less intrusive and more accurate method for scanning UDP ports on a target computer. As with TCP scanning, a UDP host discovery list of commonly used UDP ports is created.

An example of an UDP discovery list is shown in Table 4.

TABLE 4

Sample UDP Host Discovery List	
Service	Port
DNS	53
DHCP	67
BootP	69
NTP	123
NetBIOS	
File/Printer Sharing	
RPC	137
Pipes	
WINS Proxy	
SNMP	161
IKE	500

Unlike the data in traditional UDP port detection packets, the data contained within the UDP packets sent in accordance with the present invention are specifically designed to prompt a reply from the scanned host (i.e., target computer) based on knowledge of a service typically associated with the UDP port. If no information is available about the UDP port, standard data (for example, the data representing a simple ASCII character return or control character) are placed in a UDP packet. In one embodiment, an exemplary UDP data probe packet is designed to solicit a response from a NetBIOS name service that typically runs on UDP protocol at port 137. An exemplary UDP data probe for UDP port 137 is shown in Table 5. In this case, the probe is advantageously a NetBIOS node status request, which typically results in a known response from the UDP port.

TABLE 5

		Sample UDP Data Probe															
Service	Port	Data probe (hex)															
NetBIOS	137	A2	48	00	00	00	01	00	00	00	00	00	20	43	4B	41	41
		41	41	41	41	41	41	41	41	41	41	41	41	41	41	41	41
		41	41	41	41	41	41	41	41	00	00	21	00	01	---	---	---

Similar UDP data probe packets can be created for other UDP ports known to be associated with certain services, based on publicly known information in Internet Requests For Comment (RFCs), published specifications on private standards, analysis and monitoring of traffic on known ports, or well-established reverse engineering techniques. Particular protocols and particular UDP ports are subject to substantial change over time, due to changes in standards, technology, operating system, and the like. Preferably, when a data probe from a known protocol is not available for a certain UDP port, standard UDP packets with data representing at least one simple ASCII carriage return are sent.

As with TCP port scanning, multiple UDP ports can be advantageously scanned in parallel. Typically, the system sends a UDP data probe packet to a set of UDP ports on each target computer. The system waits until a UDP data reply is received from one or more target computers or until the packets sent have "timed out." If an ICMP unreachable reply is received, the hosts added to the "potentially live" list. Only those target computer IP addresses not yet determined to correspond to live or potentially live target computers after employing ICMP, TCP and UDP scanning are left on the scan list.

The three-step discovery phase, employing ICMP, TCP and UDP scanning steps, optionally is applied multiple times to increase the accuracy of identification of live target computers on the target network. In one embodiment, after an initial pass through all three steps, if there are any remaining target computer IP addresses on the scan list that have not been identified on the "live list" or on the "potentially live list," at least the ICMP and TCP steps are repeated a predetermined number of times for those remaining scan list target computers. Optionally, the UDP step may also be repeated. In particularly preferred embodiments, more intensive UDP scanning techniques, employing more ports, different data probe packets or more data packets, can be applied to the target computers remaining on the scan list to provide a more definitive list of live target computers on the network. Alternatively, traditional UDP scanning with malformed packets can be attempted.

One obstacle to the usefulness of UDP scanning is that some target computers will limit the number of ICMP responses sent within a predetermined "response latency" time period, ranging from about 250 milliseconds to about one second. Thus, if a target computer is sent twenty UDP requests to various ports in one response latency period, it may only send one or two ICMP error responses indicating that the ports are closed, even if all twenty ports tested are closed. This results in substantial ambiguity as to UDP port status. In particular, when applying traditional UDP scanning techniques, many "false positives" (i.e., falsely reported open ports) are possible.

The present invention advantageously overcomes the foregoing problem by (1) determining dynamically the latency period of the target computer, and (2) continuing the UDP port scanning of each target port for at least one entire

latency period (or until an ICMP error response or UDP data response is received) to ensure that a non-responsive UDP port is actually open and is not simply masked by the response latency. The latency period is determined by selecting a UDP port, preferably one that is known to be closed, and sending UDP datagram requests for a predetermined maximum response latency period (i.e., for a time at least as great as the target computer dead time period (approximately two seconds in particular embodiments)). The time between responsive ICMP error messages or the time between UDP response packets represents the standard latency period. This test can be repeated to confirm the latency period (dead time).

Once the latency period is determined, response validity is ensured by sending UDP requests continuously to the target port for at least the pre-determined latency time or until an ICMP error response or UDP data response is received. If an ICMP error response is received, the port associated with the prompting UDP request may be assumed to be closed. If a UDP data response is received, the associated port may be assumed to be open. If no response is received for the entire latency period, the associated port may be assumed to be open unless a router, firewall or packet loss has interfered with UDP port response.

The foregoing three-step discovery phase and optional intensive UDP scanning step are illustrated by the process flowchart in FIG. 6. As discussed above and as illustrated in FIG. 6, the discovery phase begins with a scan list 610, which is advantageously parsed into batches of IP addresses 612. In a step 620, a batch of IP addresses is obtained for processing. In a step 630, the ICMP ping test is performed as discussed above. Depending on the outcome of the test, each IP address is added to a live list 632 or remains in the present batch of IP addresses 612. As illustrated, the process also operates with a potentially live list 634 and a dead list 636.

After performing the ICMP ping test, the process determines, in a decision step 638, whether any of the IP addresses in the current batch have not been added to the live list. If no IP addresses remain, the process proceeds to a decision step 640 where the process determines whether all batches of IP addresses have been scanned. If batches of IP addresses remain, the process returns to the step 620 and obtains a new batch of IP addresses. Otherwise, the process ends.

If, in the decision step 638, the process determines that one or more IP addresses have not been added to the live list, the process proceeds to a step 650 where the TCP discovery scan described above is performed using a TCP discovery port list 652. IP addresses are added to the live list 632 or to the potentially live list 634 in accordance with the results of the scan. Then, in a decision step 654, the process determines whether any of the IP addresses in the current batch has not been added to the live list or the potentially live list. If no IP addresses remain, the process proceeds to the decision step 640 discussed above.

If, in the decision step 654, the process determines that one or more IP addresses have not been added to the live list or the potentially live list, the process proceeds to a step 660 where the intelligent UDP discovery scan described above is performed using a UDP discovery port list 662. IP addresses are added to the live list 632, the potentially live list 634 or the dead list 636 in accordance with the results of the scan. Then, in a decision step 664, the process determines whether any of the IP addresses of the current batch have not been added to the live list, the potentially live list or the dead list. If no IP addresses remain, the process proceeds to the decision step 640 discussed above.

If, in the decision step 662, the process determines that one or more IP addresses have not been added to the live list, the potentially live list or the dead list, the process proceeds to a step 670 where the intensive UDP discovery scan described above is performed using a second UDP port list 672, which advantageously includes additional ports to scan. In one embodiment, the second UDP port list 672 is the UDP service port list advantageously described below, but any combination of UDP ports may be used. IP addresses are again added to the live list 632, the potentially live list 634 and the dead list 636 in accordance with the results of the intensive scan. Then, in a decision step 674, the process determines whether the discovery testing needs to be repeated. If all IP addresses have been added to one of the three lists and no IP addresses remain, the process proceeds to the decision step 640 discussed above. If any remaining IP addresses remain that have not been added to one of the lists, then the process determines whether the scanning steps have been performed a predetermined number of times for the current batch of IP addresses. If the steps have been performed a predetermined number of times, the process proceeds to the decision step 640. Otherwise, the process returns to the ICMP ping test 630 to again process the remaining IP addresses.

As discussed above, the intensive UDP discovery scan 670 is optional. If the intensive UDP discovery scan 670 is not included, then the decision step 664 following the intelligent UDP discovery scan 660 advantageously includes the additional decision of whether to repeat the scanning process based on whether the scanning process has been performed a predetermined number of times.

#### B. Service Discovery

In one embodiment, the present invention then proceeds to examine each host (i.e., target computer) in more detail using the live list and optionally using the potentially live list as well. In service discovery, a set of common TCP service ports and a set of common UDP service ports are scanned. The TCP service discovery list and the UDP service discovery list are typically substantially larger than the TCP host discovery list and the UDP host discovery list. Each list typically includes some subset of ports that are commonly used for communications. For example, each list may include anywhere from a few to hundreds of ports. Typically, each list includes ports such as those shown in Table 6 (an exemplary list of default ports publicly reported by Microsoft® for Windows® 2000.) This list is not exhaustive, and changes in technology, protocols, network infrastructure and operating systems frequently change port requirements.

TABLE 6

Example Set of UDP and TCP ports for Service Discovery		
Service	UDP Ports	TCP Ports
NetBIOS	138, 137	
Client/Server		135
CIFS	445	139, 445
DCOM	135	135
DHCP		67, 68, 135
DNS	53	53, 139
Exchange 5.0		135
IMAP		143, 993
LDAP		389, 636
POP3		110, 995
RPC		135, 1500, 2500
SMTP		25
NNTP		119, 563
File Sharing	137	139
FTP		20, 21
HTTP		80, 443, 8000, 8080
IIS		80
IKE	500	
IRC		531
ISPMOD		1234
Kerberos	88, 464	88, 464, 543, 544, 2053
WinNT Login	137, 138	139
Macintosh File Services		548
MSN Chat/Messaging	1801, 3527	135, 569, 1801, 2101, 2103, 2105, 6665, 6667
NetMeeting		389, 522, 1503, 1720, 1731, 1723
PPTP		139
Printer Sharing	137	
SNT	162	
SQL/Pipes/RPC	137	139, 1024-5000, 1433
Telnet/Terminal		23, 3389
UNIX Printing		515
WINS	137	42, 102, 135, 137

Because there are  $2^{16}$  (65,536) possible ports on each target computer, the selection of a subset of ports is preferred because time constraints generally preclude scanning of all ports on each target computer, especially on a large network. Similarly, random selection of ports on each target computer is unlikely to be fruitful because an average target computer will run less than a dozen, and in rare cases dozens or hundreds, of services, making the probability of hitting an open port through random selection of a port inefficient and inaccurate.

In accordance with preferred embodiments of the present invention, TCP service discovery uses the large list of TCP service scan ports, a few of which are shown above, and attempts to connect to each target port on each target computer. As with Host Discovery, described above, standard TCP SYN scanning requires waiting for a SYN ACK response from each target port, or TCP "full connect" scanning requires waiting for an operating system message indicating that a three-way handshake between the target computer and scanning system has been completed. Target computers that respond positively to either of these attacks are added to a list of vulnerable computers to undergo vulnerability assessment for each open target port found.

In accordance with preferred embodiments of the present invention, UDP service discovery uses the large list of UDP service scan ports, a few of which are shown above, and attempts to connect to each target port on each target computer. As with Host Discovery, described above, the present invention may advantageously use the improved UDP data probe packets for any port recognized to be commonly associated with a particular service. Alternatively, packets that include data representing standard characters, such as one or more ASCII character returns, may be



used. Target computers that respond positively to this attack are added to a list of vulnerable computers to undergo vulnerability assessment for each open target port found. The vulnerabilities used are typically limited to those associated with the operating system of the target computer as discovered by the operating system identification system described previously, and by those vulnerabilities associated with the open ports found on the target computer. If the operating system of the target computer cannot be conclusively identified, then typically all vulnerabilities associated with the open ports found on the target computer during the service discovery system described herein are tested against the target computer.

As with Host Discovery, a more intensive UDP scanning regime can be applied to target computers that do not respond to a simple UDP data probe packet. The more intensive UDP scanning regime advantageously uses, for example, traditional UDP port scanning in combination with the optional latency time resolution described above.

After vulnerability assessment, an optional active assessment of target computers takes place. Active assessment applies knowledge of target computer operating system, open ports, as well as information recovered from the target network during vulnerability assessment, to further test known vulnerabilities against each target computer.

#### Parallel Processing of Multiple Ports

In one preferred embodiment, the present invention advantageously performs port scanning in "batches," rather than completing serial scanning of individual ports, one after another. This allows small pieces of a large scan to be completed and saved to a database, not requiring the entire results of a large scan to be held in memory all at once. For example, if a class A network is tested, hundreds of millions of ports need to be scanned merely for host discovery. The system may be advantageously adapted to technical limitations in older network equipment to scan ports of computers in a very large target network with multiple, significant IP address ranges. In the event of network or system failure, the system resumes particularly large scans following the last known successfully completed batch.

During a port scanning process, preferred embodiments of the present invention identify small portions or "batches" of the entire scan range of IP addresses, comprising, for example, 64 target computers represented in the scan range. Host discovery begins using a first batch of target computers. When all live hosts in the first batch have been discovered, service discovery of TCP and UDP services is performed on the live hosts in the first batch, along with other testing services such as operating system identification, tracerouting of network topology, and active assessment of vulnerabilities, described herein. When a batch is completed, the results are stored in a vulnerability database. Then, the next batch of 64 target computers are similarly processed, and so on, until the entire scan list has been processed for the target network.

Preferably, a large number of target computers are tested on one or a small number of ports in parallel, rather than a large number of ports being tested on a single or small number of target computer in parallel. Typically the former method prevents undue load on any single target computer, prevents network overload, and reduces the risk of triggering target network security alerts. Occasionally, the later method may be preferred if, for example, a few target computers are selectively singled out for scanning.

More particularly, in one exemplary embodiment, the present invention simultaneously uses a set of 640 packet slots. A packet slot is an advantageously allocated space in memory that keeps track of a sent packet while waiting for a response or for the packet sent to time-out. Thus, when processing a batch of target computers, up to 640 ICMP request packets can be handled simultaneously, but it will be appreciated that a different number of slots, more or less than 640, may be allocated and used simultaneously. When an ICMP response is received for one of the packets sent, or after the packet has timed-out, the slot allocated for the packet is then reallocated to send a new packet and monitor any response or time-out. Thus, packets can be kept moving at a rapid rate since there is a small average delay between sending of one packet and receipt of a response and/or time out for another packet. It is more than likely, for example, that by the time the 640<sup>th</sup> ICMP request packet has been sent, a response to one of the sent packets will have been received, thereby freeing the associated slot for handling of another packet. A similar methodology applies to the handling of TCP SYN packets, UDP host discovery, and service discovery in general. As those of ordinary skill in the art will appreciate, operating system identification and tracerouting can use similar batched parallel port monitoring. For those processes, the number of slots used is not limited to 640 slots, and can be varied for programming or scanning efficiency.

As an example, one embodiment of the present invention uses a batch size of one hundred target computers for simultaneous scanning. For host discovery, the ICMP scanning process sends 100 ICMP echo requests and monitors responses and time-outs. Because this is less than this embodiment's 640 packet slots, all ICMP request packets can be sent out, and responses or time-outs for all packets can be monitored at the same time. Assuming a TCP host discovery list of 10 TCP ports, the TCP host discovery scanner will ultimately send out 1000 packets (100 target computers×10 ports=1000 packets) to probe all listed ports of all listed computers. In this embodiment, the scanner will initially send out 640 packets in a sequential burst using all 640 packet slots. However, as discussed above, by the time the 640<sup>th</sup> packet is sent, it is probable (depending on time-out settings and how many responses the system receives) that at least one of the 640 packets sent earlier will have generated a response or a time-out, freeing that packet slot for sending one of the remaining 360 packets. By constantly monitoring the receipt of response packets and time-outs, the entire list of target computers can be tested with little wasted time.

Larger batch sizes and more packet slots can be used, but the efficiency of such use depends on how much of the target network comprises live hosts, time-out values, network response time, memory and other system resources, and the like. In general, very large batch sizes take correspondingly longer in the host discovery phase, such that by the time the system begins scanning the discovered hosts in the service scanning phase there is a greater possibility that those target computers may have changed in some way.

The above-described service discovery phase is illustrated by the process flowchart in FIG. 7. The process starts with a live list 710, which advantageously corresponds to the live list 632 generated by the host discovery process illustrated in FIG. 6. The process may also include the potentially live list 634 generated in FIG. 6; however, for ease of explanation, the potentially live list is not shown in FIG. 7. The IP addresses in the live list 710 are advantageously parsed into batches of IP addresses 712. As described below, the process

of FIG. 7 operates on a target computer vulnerability database 714 and a known vulnerability database 716.

In a first decision step 720 in FIG. 7, the process determines whether TCP SYN scanning or TCP "full connect" scanning is being performed. As discussed above, the TCP full connect scanning process requires the process to wait for an operating system message indicating that a three-way handshake between the target computer and the scanning process, the process, waits for a SYN ACK response from each target port, in which case, the target computers can be processed in parallel, as discussed above. In both cases, the process proceeds to a step 722 to obtain a first batch of IP addresses representing a first batch of live (or potentially live) target computers. If TCP full connect scanning is being performed, the process operates on a smaller number of ports on the target computers at one time. The process then proceeds to a decision step 730.

In the decision step 730, the process determines whether all the live target computers have been processed in TCP full connect scanning or whether all the batches of live target computers have been processed in TCP SYN scanning. If all the target computers or all the batches of target computers have been processed, the process ends. Otherwise, the process proceeds to a TCP service scan routine 740 wherein the process uses a TCP service discovery list 742 to identify the TCP service ports to be examined for each target computer. As described above, TCP packets are sent to the identified TCP service ports of each target computer, and the target computer vulnerability database 714 is updated for each target computer in accordance with whether a response is received or is not received from each target computer for each TCP service port scanned and using the known vulnerability database to obtain the vulnerability information for the particular TCP service ports that are determined to be open.

After performing the TCP service scan routine 740, the process proceeds to an optional UDP latency test 750 wherein the latency of each target computer is determined and stored in a latency database 752. The process proceeds from the latency test 750 or directly from the TCP service scan routine 740 to a UDP service scan 760 wherein the process uses a UDP service discovery list 762 to identify the UDP service ports to be examined for each target computer. As described above, UDP packets are sent to the identified UDP service ports of each target computer, and the target computer vulnerability database 714 is updated for each target computer in accordance with whether a response is received or is not received from each target computer for each UDP service port scanned and using the known vulnerability database to obtain the vulnerability information for the particular UDP service ports that are determined to be open.

After completing the UDP service scan routine 760, the process proceeds to a decision step 770 wherein the process determines whether a response has been received from all the scanned service ports. If not, the process determines whether the scanning of the current target computer or batch of target computers has been performed a predetermined number of times. If all the scanned service ports have responded or if the scanning has been performed a predetermined number of times, the process returns to the step 720. Otherwise, the process returns to the TCP service scan routine 740 to scan the service ports that have not responded.

#### C. Banner Grabbing

"Banner grabbing" is a method of stripping relevant information from packets received from a target computer.

In one embodiment, for each target computer in the live list, an attempt is made to perform banner grabbing against each open TCP port and each open UDP port discovered during the service scan phase. If information is successfully obtained from a responsive TCP or UDP packet, this information is stored in a database in association with the target computer from which the information is obtained. The stored information is subsequently used to collect vulnerability information on the target computer, and the vulnerability information is also stored in a database.

For each open TCP port located during the service discovery phase, a TCP data probe is sent to that port if the port is known to be typically associated with a particular service. Thus, for example, when sending a TCP data probe to a target TCP port 80 on a target computer, where TCP port 80 is well known as the common service port for HTTP, the system sends an HTTP-compliant GET request, and strips useful information from any TCP response packet sent back from the target computer. As noted in the above tables, similar commands can be used for other ports with well known services running thereon. The useful information obtained may advantageously be transformed or translated into a readable or processable form (e.g., text) and stored for later reporting.

For each UDP port discovered during the service discovery phase, a similar UDP data probe is sent to each UDP port on the target computer known to be typically associated with a service. Thus, for example, from the foregoing tables, UDP port 137 is known to be associated with a NetBIOS service. In one embodiment, the system sends a NetBIOS Node Status Command, for example, and strips and stores useful information from any response packet sent by the target computer. Again, the information obtained in this manner may advantageously be transformed or translated into a readable or processable form (e.g., text) and stored for later reporting.

#### IV. Tracerouting

In one embodiment, for each target computer in the live list, an attempt is made to perform an ICMP traceroute between the system and the target computer. If the ICMP traceroute is incomplete, an attempt is made to perform a TCP traceroute between the system and the target computer. Based on the traceroute results for each target computer, a map of the target network topology is created and stored in a database.

Initially, traditional ICMP tracerouting is performed. A number of packets are sent to the target computer with varying TTL (time to live) values in the TCP/IP header, starting with a TTL of 1. If the ICMP echo request packet fails to reach the destination target computer, it will return an ICMP destination unreachable packet containing the IP address of the router/host that the packet was returned from. Packets that arrive at the target computer itself return an ICMP echo reply packet. Based on the known TTL value, the number of "hops" between the system and the target computer, and the router/host at each "hop" can be mapped out. When all packets have arrived or timed out, if there is a router/host and IP address associated with each "hop" between the system and the target computer, the traceroute step is completed for that target computer. The ICMP traceroute is continued for a number of pre-determined passes until a complete trace is received or the maximum number of passes is completed. If, after the pre-determined number of passes the tracerouting of the target computer is incomplete, then a TCP traceroute is attempted.

TCP tracerouting works similarly to ICMP tracerouting, except TCP SYN packets are used. In some instances, TCP packets are more reliable at completing otherwise incomplete traces, because, for example, ICMP packets are sometimes blocked by firewalls or routers. Typically, the TCP port chosen is taken from the list of open ports discovered for the target computer that is the subject of the traceroute. As with ICMP tracerouting, the TTL value in each SYN packet is incrementally increased, and the return of ICMP unreachable packets and SYN ACK packets is monitored for all "hops" between the scanning system and the target host. Through the combination of ICMP tracerouting and TCP tracerouting, a complete map to each target computer, and collectively a relatively complete map of the target network topology, is advantageously created.

The network maps are an attempt to represent in the most concise way possible the topology of the network as discovered by tracerouting methods. Basically, packets are sent to each discovered host in the scanned network with the Time-To-Live field set to the hopcount at which a packet indicating failure to transmit due to that hopcount will be sent by the machine that is that many hops away from the scanning machine. By increasing the TTL and storing the IPs of the machines that respond until getting a response from the host in question, a string of IPs is built up that represents the machines that led up to that host. The procedure is far from perfect. A machine at a given hop may not respond within the time expected by the tracerouting algorithm, or it may not respond at all if it is a firewall that blocks the traceroute packets. Any approach that can exist must deal with this uncertainty in some way. In one embodiment, any unknowns are assumed to be firewalls.

The algorithm typically employed is presented with a set of discovered hosts and each host includes a IP address that represent what the tracerouting algorithm initially determines are the machines leading up to them. This string may include unknowns which are denoted by 0xFFFFFFFF or 255.255.255.255. The algorithm operates as follows:

**CONDENSE UNKNOWN:** Condense consecutive unknowns into a single hop with unknown IP address (0xFFFFFFFF). If a traceroute is performed to a given hop and an unknown is received, it is probable that every hop after that will also be unknown because the first unknown hop is a firewall blocking the packets.

**UNKNOWN RESOLUTION:** Attempt to resolve unknowns by looking through the other routes to see if a route exists with the IP addresses on either side of the unknown in this route that are connected by all known machines. If found, replace the unknown with those machines. This attempts to eliminate spurious unknowns so they are not mislabeled as firewalls. It will not eliminate actual firewalls as they will typically be unknown.

**BUILD SET OF NODES:** Build up a list of structures (structs) of type ROUTENODE. This is a list in which the IP addresses are guaranteed unique (except for unknowns as detailed later) and represents the set of all IP addresses discovered in the network (i.e., the hosts discovered by the scanning process and the hops leading up to them discovered as by the tracerouting process.) The m\_pHost member of each ROUTENODE struct will be NULL if the IP address is not one of the hosts discovered in the scanning process and will point to the host data discovered by that process otherwise. As mentioned, the known machines are unique by IP address. The unknowns are further qualified by the IP address immediately preceding them. As a result, unknowns with the same IP address

preceding them are typically considered to be the same firewall and are therefore represented by a single ROUTENODE in the list.

**FILL IN CONNECTIONS:** Each ROUTENODE struct has a list of ROUTENODE pointers in it which represent the machines that ROUTENODE is directly connected to. The algorithm next fills this in to guarantee that the IP addresses are unique in each list of connections. Note that these pointers point to other ROUTENODE structs in the list (i.e., the list is typically a self contained bunch of data.)

**COMPUTE DISTANCE FROM INTERNET:** The traceroute information is once again traversed for each host and for each IP address. The corresponding ROUTENODE struct is looked up in the list and the struct's m\_nDistFromInternet is set according to the hopcount in the route. This is typically done to represent a maximum hopcount. In other words, the m\_nDistFromInternet fields are initialized to 0, and then, if the hopcount in a traceroute IP string is greater, the m\_nDistFromInternet is set to that hopcount. For example, an IP address (let it be called "d," for example) may appear in one route as a-b-d and in some other route as a-b-c-d, where "a," "b," and "c" are other IP addresses on the route to the IP address c. When this part of the algorithm completes, d will have m\_nDistFromInternet=4. This information is used by the algorithm to know if one machine precedes another. Note that this maximum of discovered hop lengths is a heuristic for the sake of making the problem reasonably computable. It could go awry if a particular machine should be drawn as connected to another machine (because it is directly connected to it) and precedes it in one route, but it is not drawn that way because the machine is in some other route that bumps its m\_nDistFromInternet to the same or more than that of the machine it is connected to in the other route. This situation is rare, and this heuristic is typically acceptable.

**BUILD ROUTER LIST:** The algorithm then traverses the ROUTENODE list in order to build up a list of ROUTER structs. A ROUTER struct contains a pointer to a ROUTENODE believed to be a ROUTER and a list of ROUTENODE pointers that are leaf (end of route) machines it is connected to. A leaf is any ROUTENODE that is directly connected to only one machine (besides itself.) For the purposes of this algorithm, a ROUTER is that machine. This stage of the algorithm builds up a list of ROUTER structs, and, within each ROUTER, builds a list of leaf ROUTENODEs connected to that ROUTER. The algorithm also marks a Boolean field within the ROUTENODEs that are being designated ROUTERs as it will become necessary to know which machines in the ROUTENODE list are routers without wanting to consult the ROUTER list.

Next the refinements to the algorithm are discussed, but first it is useful to discuss how the map would be drawn if the algorithm finished here. At this point, the algorithm has a set of ROUTERs, and within each of those, has the set of leaves connected to that ROUTER. Prior to the refinements about to be detailed, the map renderer would draw each ROUTER as a sphere and all of its leaves as small spheres embedded in a ring in orbit around the ROUTER sphere. The algorithm would then look for other ROUTERs whose m\_nDistFromInternet fields are less than this ROUTER's m\_nDistFromInternet (i.e., look for a ROUTER that is directly connected to this ROUTER and precedes it.) The algorithm takes the greatest m\_nDistFromInternet router that it finds and connects this ROUTER to it visually. If it

finds no preceding ROUTERS, it connects this ROUTER to the internet cloud in the center of the map. The refined algorithm still does all that is described above, but typically not before doing a bit more processing.

ROUTER PROMOTION: There are circumstances when the algorithm as it has been detailed thus far would produce a map in which a ROUTER has leaf nodes surrounding it all of which have the same subnet, for example they might all be 149.17.1.xxx. The first 1, 2, or 3 digits (8 bits each) will be the same. That ROUTER has a ROUTER of less m\_nDistFromInternet directly connected to it that has exactly 1 leaf node and that leaf node has the same IP address digits in common as all the ones in the first ROUTER's leaf list. In this case, although the traceroute data suggests that the first ROUTER and that single leaf are different machines, they are likely two NICs (Network Interface Cards) on the same machine. Hence, the algorithm "promotes" that ROUTENODE by adding its IP address to the m\_nOtherNics member of the first ROUTER's ROUTENODE and then removing it from the other ROUTER's list of leaves and marking that ROUTER's ROUTENODE to be no longer a ROUTER. That ROUTER is also removed from the ROUTER list. On the map then, the ROUTER's primary IP address is represented as usual but now there is a grey bar and the other IP address underneath it. The algorithm supports doing multi-homed ROUTERS and will represent all the IPs that get promoted. The multi-homed case is easily imagined by extending the dual-homed example discussed here.

OTHER TYPE OF ROUTER PROMOTION: The preferred algorithm performs another type of router promotion as well. If exactly one leaf around some ROUTER fulfills a heuristic, it will be promoted to another NIC on the ROUTER. In order to test for the heuristic, the machine needs to be a host discovered during the scanning process because that is where the information comes from. It is noteworthy that most leaves will be discovered hosts. Although the algorithm does not demand this, and traceroute information may produce non-discovered-host hops that only have one thing connected to them and thereby make leaves out of them, since tracerouting is being done to the discovered hosts, the leaves will tend to be discovered hosts. The heuristic is, in one embodiment, as follows: if TCP 2001, 4001, 6001, or 9001 is discovered, or if UDP 520, or both UDP 23 and 79 are discovered, or if the discovered operating system is Cisco, that ROUTENODE is assumed to be a router-IP address and it is promoted (IP address added to the ROUTER's m\_nOtherNics and ROUTENODE pointer removed from its leaves).

FIREWALL PROMOTION: Firewall promotion is similar to the router promotion heuristic detailed above. In one embodiment, if exactly one leaf around a firewall satisfies the heuristic then it gets promoted into the firewall's other-NIC list resulting in the map printing "Firewall," but now the known IP address is underneath it separated by the gray bar. The heuristic, in one embodiment, is if TCP 256, 257, 264, 265 is found on a machine then it is assumed to be a firewall.

NUMBER FIREWALLS: This is the last step of the algorithm. The firewalls are sorted in order of the IP address that precedes them and are then numbered so that the map can print "Firewall-1, Firewall-2, . . . etc." on the report. In this manner, an entire network map can be relatively accurately stored internally, and translated into a visual representation for reporting.

## V. Vulnerability Identification and Active Assessment

For each known TCP port, UDP port, and operating system, the known vulnerabilities for that configuration are stored in a vulnerability database based on vulnerability identification codes. Advantageously, for many vulnerabilities in the vulnerability identification database, a methodology for testing the vulnerability can be written into an automatic script, which will assess the actual weakness of the target system to the suspected vulnerability. In one embodiment, these scripts are prepared in a assessment security scripting language, and are preferably prepared in FASL. FASL is a scripting language based on C++ and Java implementation, in one embodiment. FASL provides an adjustable, automated language for security testing various vulnerabilities. Multiple FASL scripts advantageously may be run in parallel. For example, in one embodiment up to eight simultaneous scripts may be run. Each FASL script, for example, will respond with a success or failure result noting whether the target computer was vulnerable or not to the given vulnerability identification code. The information collected by the FASL script from the target computer, and the success or failure of the attempt, are stored in a database related to the target computer for later reporting, for use in additional vulnerability testing, or for repeating additional testing.

## VI. FASL Scripting Language

The implementation of the FASL language is similar in construction to C++. In one embodiment, FASL includes member functions in structure objects, constructors and destructors in structure objects, inheritance in structure objects, arrays of scalar types, new scalar type "binary" and associated functions, string constants that support embedded hex codes including embedded zero bytes, RPCCheck() and SMBCheck() functions for RPC and Netbios checks, debugMessage() on all scalar types that produces hex output on binary type, recursion, function overloading, reference parameters, and support for Active-Assessment.

In one particular implementation of FASL, all variables are stack variables. In addition to permitting recursion, this also allows a parse tree to be built once and then used to execute the same script by multiple threads. The execution architecture in FASL treats the parse tree as a read-only data structure. Hence, other than considerations such as two instances of the same script attempting to instantiate a Socket on the same target computer IP and port (which should not happen in practice), FASL is completely thread-safe.

### A. Scalar Data Types

Scalar data types are those native types in the language that take up one slot on the execution stack. Basically, any variable that can be declared is a scalar type, an array of a scalar type, or an instance of a structure. The scalar types are as follows:

TABLE 7

FASL Data Types	
Scalar Type	Definition
void	Function return type
int	64 bit signed integer
string	String of printable characters of arbitrary length, terminated by a null
Binary	String of arbitrary bytes of arbitrary length, terminated by a null. Type keeps track of its length

TABLE 7-continued

FASL Data Types	
Scalar Type	Definition
Char	8 bit signed integer (interchangeable with int)
Boolean	True or false (not interchangeable with int)

Typically, a string is NULL terminated. This means that internally the string is represented as a string of characters with a zero marking the end. This zero's existence is always implicit, however. In other words, there is no string function that will return you the zero. The length of the string is computed by counting the number of characters up to but not including the zero. This also makes it so that if a constant is assigned to a string of, for example, "this is a string\x0 with stuff beyond the zero." the value of that string will be "this is a string" and its length will be 16. Type binary typically does not use a zero or any other delimiting character in its internal representation. The string constant example with the embedded zero would have everything after the zero also in the binary.

For example,

```

binary b = "1234";           // length = 4.
string s = "1234";           // length = 4.
binary b = "zzzz\x0ssss";    // value = "zzzz\x0ssss", length = 9.
string s = "zzzz\x0ssss";    // value = "zzzz", length = 4.

```

Types which are equivalent to int (interchangeable in expressions with int) are as follows:

TABLE 8

Additional int-like FASL types	
Type	Description
Attack	specifying type of attack
Operatingsystem	specifying target OS of a script
Protocol	specifying protocol used by a Socket
Returnvalue	specifying status of script (script return value)
Ipservice	IP type of service
Ipoptions	IP options
Ippoffset	IP offset

Keywords that indicate a constant value of a given type are as follows:

TABLE 9

Additional FASL Constants	
Constant Keyword	Value
null	0 (int, char, boolean)
true	1 (boolean)
false	0 (boolean)

The types above that are equivalent to int also have keywords indicating constant values of those types. They are (respective to above) as follows, for example:

TABLE 10

FASL Type Constants	
Type	Constants
attack	ATTACK_UNKNOWN, ATTACK_INTRUSIVE, ATTACK_DOS, ATTACK_NONINTRUSIVE
operatingsystem	OS_UNIX, OS_MAC, OS_WINDOWS, OS_UNKNOWN, OS_ROUTER
protocol	TCP, UDP, ICMP, IGMP
returnvalue	RETURN_SUCCESS, RETURN_FAILURE, RETURN_TIMEOUT
ipservice	LOWDELAY, THROUGHPUT, RELIABILITY, MINCOST, CONEXP, ECTCAP
ipoptions	EOL, NOP, RR, TS, SECURITY, LSRR, SATID, SSR, RA
ipoffset	RESERVED, DONT, MORE

### B. Statements

A FASL script is typically a list of statements. These statements are generally separated by semicolons. The exception to this rule is that statements that involve blocks (other statement lists enclosed by {and}) do not typically require semicolons to separate them. A semicolon constitutes a valid (empty) statement so it does not hurt anything to put semicolons after statements involving blocks but it accomplishes nothing. A new block represents a new scope. Any variables declared in that scope will be accessible only in that scope, and if they have destructors, those destructors will be called at the end of the scope. Variables can be declared with the same name as variables in scopes that enclose the scope in which they are declared. They will override the variables in the outer scopes which are otherwise accessible to the enclosed scope. Variables of the same name as other variables in the same scope will generate an error however. A statement can be one of the following, each of which will be detailed further below:

Function declaration: A named list of statements with parameters that appear to the list of statements as declared variables in their scope but whose values are copied from expressions or else refer to variables supplied by a function-call which is a type of expression. Function declarations may only occur in the topmost scope of the script.

Structure declaration: A declaration of an entity which can have both data members and member functions. The member functions all have an implicit first parameter which is a reference to the object instance upon which the function is being called.

Variable declaration: A declaration of a variable which then exists in the scope in which it was declared. Variables can be simply declared or as part of their declaration they can be initialized by assigning them to an expression or constructed by supplying constructor parameters in parenthesis after the variable. Note that scalar types have no constructors. On types that do have constructors, the parameter list must match a constructor that exists for that type.

Expression: This can be a function call, a constant value, a variable, a member selection (which is either a member variable or member function), and arithmetic and logical operations on variables, constants, etc.

While loop: This is a control structure that has a block/statementlist/scope that is executed for as long as a given condition (expression) resolves to true.

Repeat loop: This is a control structure that executes its block a given integer amount of times that an expression resolves to.

If statement: This is a control structure that executes a block upon an expression evaluating to true and executes an optional else block if it is false.

For loop: This control structure has 3 lists of expressions separated by semicolons enclosed in parentheses. Within the lists (which can be empty, have one expression, or more than one) if there are more than one expression, they are separated by commas. The first list is executed initially. The next is a condition that must be true for the block to execute (in a list of expressions all must be true), and the last is the iterator list of expressions that gets executed each time through the loop.

Block: Note that many statements own blocks. A block may also exist on its own without being owned by a control statement or function declaration. An unowned block may be coded in circumstances when construction and destruction of something needs to happen at a given point that cannot be accomplished in any of the "normal" scopes.

#### C. Function Declarations

Functions are declared using the "function" keyword. They are declared as follows:

```
function [

```

When a return type is not specified, int is implicitly assumed. The argument-list can be empty (i.e., "function <functionname>() . . ."). Every script typically needs a "function faslmain()" and will error if one is not found. The function faslmain() can be explicitly called from anywhere in the script but if no call to faslmain() is found in the topmost scope, a call is implicitly inserted at the end of the script. Functions may be overloaded, such that multiple functions with the same functionname can exist provided their argument-lists are different. The argument-list has the form "<argument>, <argument>, . . . , <argument>". An argument can be the following forms:

Pass by copy. <scalar-type><distinct-variable-name>: Anything the function does to this variable is not reflected in the corresponding expression supplied in the call. The expression can be constant.

Pass by reference. <scalar-type>&<distinct-variable-name>: Anything the function does to this variable is reflected in the corresponding variable supplied in the call. The call must supply an expression that reduces to a variable corresponding to this argument (i.e., only variables or structure member selections that reduce to variables.)

The (possibly empty) <body> consists of a list of statements. If a statement which is "return <expression>" is encountered anywhere in this body, the execution of the function ceases at that point and control is returned to the caller of the function with the value returned being the value that function presents to the expression in which it participates. Functions can be called without using the value they return as well.

For example,

```
function faslmain()
{
    int x;
    x = 5;
```

-continued

```
}
function string DoStringStuff(int x)
{
    return intToString(x);
}
function void DoStringThing(string& s)
{
    s = "the string";
    // DoStringThing(szThing) will set
    // szThing to "the string"
}
```

#### Variable Declarations

Variable declarations are declared as follows:

<typename><var>[,<var> . . .];

In the case of structures, <typename> is typically expressed as "structure <structurename>" or, alternatively, a structure variable of type "OBJECT" can be declared as "OBJECT o;".

A <var> is one of the following:

<identifier>: This is a simple declaration. If <typename> is a structure and the default (no arguments) constructor exists, it is called.

<identifier>=<initializer expression>: The expression is evaluated and the result assigned to the variable at declaration time. Note that the type returned by the expression must be compatible with <typename>.

<identifier>(<params>): <params> is a comma separated list of expressions. The constructor matching that signature in <typename> is called. Note that <typename> needs to be a structure.

<identifier>[<array-dimension-expression>]: <typename> must indicate a scalar (non-structure) type. The expression must resolve to a constant integer (i.e., no variables or function calls in it). This declares <identifier> to be an array variable of given dimension of <typename>.

For example,

```
OBJECT o;
structure OBJECT o;
OBJECT o(5, 6, "something");
int x = 8, y, z(6 + 7);
```

#### D. Structure Declarations

Structures are declared using the "structure" keyword. They are declared as follows:

```
structure <structurename> [extends <parentstructurename>]
{
    <member-list>
};
```

The parameter <member-list> is a possibly empty list of <member>s. A <member> is typically either a function declaration or a variable declaration. The only syntactic difference between these things when used in structure declarations versus not is that member variables cannot have initializer expressions or constructor parameters. Member variables can be arrays however.

When "extends" is used, this structure "inherits" all the members of the parent. Any member functions in this

structure that are the same name and signature as one that is in the parent structure will override that parent structure's function.

A "constructor" is a member function that gets called when a structure variable is declared and a "destructor" is a member function that gets called when a variable goes out of scope.

Any member function whose name is <structurename> is a constructor and can have any or no parameters. There can be zero or one member function whose name is ~<structurename> and has no parameters, and this is the destructor. Some subtleties that are not immediately apparent are as follows: If this structure extends an object that has constructors and/or a destructor or has member variables that are structures that have constructors and/or a destructor, then each constructor in this object will implicitly contain a call to the default constructor in the parent and/or owned object(s) and similarly the destructor for this object will have call(s) to destructors of parent/owned object(s). Further, if a parent or owned object has a constructor and/or destructor and this object does not, one will be created in this object for the purpose of calling all the constructors/destructors of parent/owned objects. Constructors and destructors should be declared as "function <structurename> . . ." or "function ~<structurename> . . ." (i.e., no return type.)

Usage of declared structure variables is accomplished with member selection using the "." character. For example:

---

```
OBJECT o;
o.m_intmember = 5;
o.m_ownedstructuremember.m_x = 8;
o.DoSomething(5, 6, "stuff");
```

---

### E. Expressions

Expressions are variables, function calls, constants, member-selections, and combinations of these things with operators serving as the connectors. For the purposes of this discussion, expressions may be defined recursively, such that, for example, <expression>+<expression> is an expression. In expressions involving operators (assignment, arithmetic, logical, etc.) it is customary to refer to the expression on the left as the "lvalue" and the expression on the right as the "rvalue." In FASL as most languages, the lvalue and rvalue in any expression must return compatible types (usually the types must be the same but sometimes, as with int and char in FASL, different types can be compatible). Additionally, some operators only work on some types. There is also the notion of operator precedence, meaning that in the absence of overriding parentheses, some operators will be given precedence over others. So for example, 3-5+7 will be evaluated left to right but 3-5\*7 will multiply 5 by 7 first and subtract the result from 3. A list of all the operators in order of precedence follows:

### Functions and Constants

function calls: <function-name>([<argument-list>]); <argument-list> is a possibly empty, comma-separated list of expressions. There must be a function somewhere in the script which takes parameters that match the return types of these expressions in the order in which they occur. In the case of functions that take references, the corresponding expression in the <argument-list> of the function call must resolve to a variable (not a constant or arithmetic/logical/relational expression of multiple variables). The type of a function call expression is the return type of its corresponding function declaration.

string constants: "sample string\x3f"; these are string values enclosed in quotes which mostly represent the literal characters that comprise the string. The exceptions to this are the escape characters: "\t" is tab, "\n" is newline, "\r" is carriage return, and "\\x[0-9a-fA-F][0-9a-fA-F]\*" are embedded hex codes that resolve down to one character. Note the regular expression simply means for example "\x3F-sample" will resolve down to a 47 (\x3F) followed by a '-' followed by a 's,' etc. Any embedded zeros ("\x0") will terminate the string when the constant is used in string expressions, but when used in binary expressions, the total string will get processed to resolve the escape sequences and then converted to type binary. String constants typically have the same precedence as functions.

char constants: For example, 'A'. These are treated the same as string constants, except they are single quotes and single character (which may be represented with escape sequences).

int constants: These can be decimal digits as in "1048576" or a hexadecimal number as in "0x100F."

### Unitary Operators

++<variable>: increment value of variable and then return its value. Works on int and char.

--<variable>: decrement value of variable and then return its value. Works on int and char.

<variable>+: return value of variable then increment it. Works on int and char.

<variable>--: return value of variable then decrement it. Works on int and char.

~<expression>: negate whatever is returned by <expression> and return that value. Works on int and char.

~<expression>: flip bits of whatever is returned by <expression> and return that value. Works on int and char.

!<expression>: logical negation of <expression>. Works on boolean.

sizeof(<typename> or <variable>): returns int that is how many stack cells is occupied by the <typename> or <variable>.

### Member Select Operator

<structurevariable>.<member>[.<member> . . .]: returns whatever type/value the rightmost <member> is.

### Power Operator

<expression> power <expression>: lvalue and rvalue must be int or char. If int and char mixed, promote char to int.

### Multiply Operators

<expression>\*<expression>: multiply . . . lvalue and rvalue must be int or char. If int and char are mixed, promote char to int.

<expression>/<expression>: divide . . . lvalue and rvalue must be int or char. If int and char are mixed, promote char to int.

<expression>%<expression>: modulo . . . lvalue and rvalue must be int or char. If int and char are mixed, promote char to int.

### Addition Operators

<expression>+<expression>: add . . . same type rules as multiply operators but also does concatenation of two strings (returns string) or two binaries (returns binary). If one of the expressions is a string constant and the other is a binary variable, the string constant will be considered a binary constant.

<expression>-<expression>: subtract . . . same type rules as multiply operators but also works on strings (not binaries

as operator+does). In the case of strings, subtraction removes all substrings from lvalue that match rvalue and returns the resulting string.

#### Bitwise Operators

<expression><<<expression>: shift left one bit (effectively multiply by 2). Same type rules as multiply operators.

<expression>>><expression>: shift right one bit (effectively divide by 2 losing remainders). Same type rules as multiply operators. Most significant bit (sign bit) is repeated in next bit over.

<expression>&<expression>: bitwise and. Same type rules as multiply operators.

<expression>|<expression>: bitwise or. Same type rules as multiply operators.

<expression>^<expression>: bitwise exclusive or. Same type rules as multiply operators.

#### Relational Operators

All of these operators typically return boolean regardless of the types on which they operate (in contrast to most operators which return the same type as that on which they operate). Unless otherwise noted, lvalue and rvalue can be int/char and int/char or string and string. If they are strings then the comparison is alphabetic case sensitive.

<expression><<<expression>: Return true if lvalue is less than rvalue.

<expression><=<expression>: Return true if lvalue is less than or equal to rvalue.

<expression>>><expression>: Return true if lvalue is greater than rvalue.

<expression>>=<expression>: Return true if lvalue is greater than or equal to rvalue.

<expression>=<expression>: Return true if lvalue equals rvalue.

<expression>!=<expression>: Return true if lvalue does not equal rvalue.

<expression>in<expression>: This only works on string and binary expressions. Return true if lvalue occurs as a substring/subbinary pattern in rvalue.

#### Logical Operators

These operators expect lvalue and rvalue to return boolean and the operators return boolean.

<expression>||<expression>: "or else" . . . if lvalue is true then return true without evaluating rvalue, else return whatever rvalue returns.

<expression>&&<expression>: "and then" . . . if lvalue is false then return false without evaluating rvalue, else return whatever rvalue returns.

#### Assignment Operators

These operators insist on type compatibility which means equality of types most of the time. Exceptions are: when mixing ints and chars, lvalue type prevails, and when assigning string constants into binary variables the string constant becomes a binary constant. lvalue and rvalue must resolve to scalar types and lvalue must resolve to a variable.

<expression>=<expression>: Simple assignment: copy rvalue into lvalue.

<expression>\*=<expression>:

<expression>/=<expression>:

<expression>%=<expression>:

<expression>+<expression>:

<expression>-<expression>:

<expression>>=<expression>:

<expression><<=<expression>:

<expression>&=<expression>:

<expression>|<expression>:

<expression>^=<expression>: All the above perform the operator preceding the "=" on the rvalue according to the rules of that operator specified above and then put the result in the lvalue.

#### Conditional Expression

There is also a construct identical in principle and syntax to the conditional expression in C/C++. Its syntax is:

```
(<expression>)?<expression-eval-if-true>:<expression-eval-if-false>;
```

If the expression in parentheses evaluates to true then the expression after the question mark is executed otherwise the expression after the colon is executed. The expression in parentheses must resolve to boolean and the other two expressions must be type compatible. The return type/value of a conditional expression is the return type/value of the expression that executes after evaluating the conditional expression.

#### F. Control Structures

##### While loops

A while loop is expressed as:

```
while(<expression>)
{
    <statement-list>
}
```

or,

```
while(<expression>)
    <statement>;
```

This evaluates <expression> which must return boolean and executes the <statement-list> or <statement> and then reevaluates <expression> for as long as <expression> returns true. If <expression> returns false the first time then the <statement-list> or <statement> is never executed.

##### Repeat loops

```
repeat(<expression>)
{
    <statement-list>
}
```

or,

```
repeat(<expression>)
    <statement>;
```

This evaluates <expression> once (in contrast to most loop constructs which execute their "iterator" expression each time before executing the block), and it must return an int or char indicating how many times to execute the block. The block is then executed that many times. In case the implications of this are not clear, note that in a while loop, for example, the body of the loop must do something that eventually causes the <expression> to evaluate to false, otherwise the loop goes on forever. But if the repeat loop's block does something that would cause the <expression> to evaluate differently, this makes no difference as the <expression> is only evaluated once before the block is executed. So "v=2; repeat (v) {v=5;}" will not cause an infinite loop. Rather, it will assign 5 to v twice.



```

    If statements
    if (<expression>)
    {
        <statement-list>
    }
    or,
    if (<expression>)
    {
        <statement-list>
    }
    else
    {
        <statement-list>
    }

```

For the sake of brevity, all possible options of using a single <statement> instead of {<statement-list>} have not been enumerated, but those options exist. If <expression>, which must resolve to boolean type, evaluates to true then the first block is executed. If it evaluates to false and the else block exists, it is executed, otherwise no block is executed.

```

    For loops
    for (<expression-list>; <expression-list>; <expression-list>)
    {
        <statement-list>
    }
    or,
    for (<expression-list>; <expression-list>; <expression-list>)
    {
        <statement>;
    }

```

The <expression-list> can be nothing, it can be a single <expression>, or it can be a comma-separated list of expressions. The first <expression-list> gets called initially and can be any mix of expressions as the return values are ignored. The middle <expression-list> is the condition that determines whether the block will be executed and whether the loop will continue to be iterated. All expressions in this list must return boolean and they all must return true in order for the block to be executed. The last <expression-list> can be any mix of expressions and it gets executed after each time the block has been executed.

The classic example is "for (x=0; x<8; x++) {<statement-list>}" This will set x to 0, test whether x is less than 8, find that to be true and execute the block, increment x to 1, test whether x is less than 8, find that to be true and execute the block . . . etc. until x is 8 which will cause the condition to be false and loop execution terminates.

#### G. Native Objects

The grammatical/syntactic elements of one embodiment of the FASL language have been specified above. Notice that these things are sufficient to perform calculations and implement algorithms, but there is nothing about them that allows for things like sending and receiving data over the network and there are certain operations on the scalar types that would be nice to have but are not expressible with the grammar thus far specified. For this reason there exist structures and functions declared in the scope that owns the script scope (a "meta-scope" if you will) that allow access to this functionality. As used herein "Faultline" refers to the overall network security system.

#### FASL Object

Every script has access to a variable named "FASL" whose type is "FASL\_OBJECT." This variable's type is specified as follows:

```

struct FASL_OBJECT
{
    private:
        string      m_szName,
                   m_szDescription,
                   m_szSummary,
                   m_szReturnString;
        RETURNVALUE m_eReturnCode;
        int         m_nDefaultPort,
                   m_nIPAddress;
        ATTACK      m_eAttackType;
        OPERATINGSYSTEM
                   m_eOSMajor,
                   m_eOSMinor;
        PROTOCOL    m_eDefaultProtocol;
    public:
        function FASL_OBJECT()
        {
            m_nDefaultPort = 0;
            m_nIPAddress = 0;
            m_szReturnString = "Return string not set."
        }
        function void setScriptName(string s)
        {
            m_szName = s;
        }
        function void setScriptVulnerabilityCode (int
nFaultlineID)
        {
            // This sets the vulnerability id as it exists in
the Faultline database that uniquely
            // identifies the vulnerability being sought by
the script.
            m_nFaultlineID = nFaultlineID;
        }
        function void setScriptDesc(string s)
        {
            m_szDescription = s;
        }
        function void setScriptSummary(string s)
        {
            m_szSummary = s;
        }
        function void setScriptAttackType(ATTACK e)
        {
            m_eAttackType = e;
        }
        function void setScriptReturn(string szReturnString,
RETURNVALUE eReturnCode)
        {
            m_szReturnString = szReturnString;
            m_eReturnCode = eReturnCode;
        }
        function void addValidPort(int n)
        {
            // When a script successfully compiles, it will
execute all these that it finds
            // in the main scope. It builds a list of ports
on which to run the script.
            m_nValidPort = n;
        }
        function void setScriptPort(int n)
        {
            m_nDefaultPort = n;
        }
        function void setScriptOS(OPERATINGSYSTEM eMajor,
OPERATINGSYSTEM eMinor)
        {
            // When a script successfully compiles, it will
execute all these that it finds
            // in the main scope. It will use this
information to decide whether this script
            // needs to be run on a given host.
            m_eOSMajor = eMajor;
            m_eOSMinor = eMinor;
        }
        function void setScriptProtocol(PROTOCOL e)
        {
            m_eDefaultProtocol = e;
        }
}

```

-continued

```

    }
    function int getIPAddress( )
    {
        return m_nIPAddress;
    }
    function int getScriptPort( )
    {
        return m_nDefaultPort;
    }
    function string strTok(string& s, string szDelimiters)
    {
        // Like strtok in UNIX, this skips past any
        // characters at the beginning of the string
        // that are in szDelimiters. Return the substring
        // that includes all the leftmost
        // characters up to but not including the next
        // instance of a character in
        // szDelimiters, or the end of the string
        // whichever comes first. Remove the
        // reference).
        // returned string from s (note it is a
        // reference).
        return STRTOK(s, szDelimiters);
    }
};

```

When a script is successfully compiled, all statements of the form "FASL.memberfunction" that are in the main scope are executed (and no other statements in the script are executed at this point). The information that these member functions set goes into the script object's data which allows the system to make some intelligent decisions on when and how to run the script. FASL.addValidPort(nPort) can be called any number of times and will result in a list of unique ports being built up. When the system is run, it will either find that FASL.addValidPort(nPort) was not called in the script's main scope in which case the script will be run once per host and FASL.getScriptPort() will return 0 within that script. If FASL.addValidPort(nPort) does occur in the main scope of the script, the system will execute the script once for each unique port on a given host, and FASL.getScriptPort() will return whatever port the system passes in for that execution. FASL.setScriptOS( ) operates along a similar principle (i.e., by calling it you are requesting that the script be executed only on hosts whose OS has been specifically determined whereas not calling it implies that the script is to be called regardless of OS). Note that calling this multiple times does not make a list of OSes like the addValidPort makes a list of ports. The last call to setScriptOS is the one that gets used.

Upon entry into the script's scope, the m\_nIPAddress member of the FASL variable has been set to the IP address of the target machine upon which the FASL script is to run. All network activity that takes place in a FASL script uses this IP address as its destination, hence in the functions and objects about to be specified, the IP address is never a parameter. Calling FASL.setScriptReturn(string, RETURN-VALUE) sets the m\_szReturnString member which is then printed by the script executor command line application "fasl.exe" upon completion of executing the script. The other calls, most notably setScriptOS( ), set info that is used by Faultline.exe to determine whether or not to run the script on a given host. This is to say, the system will execute only the "FASL.xxx( )" statements in the script, then the system examines the memberdata and decides whether or not to run the whole script on that IP address. In one embodiment, the constructor of FASL\_OBJECT is actually not called on the FASL variable. Thus, any variable other than m\_nIPAddress is sometimes not be initially set to anything. Structures may

be derived from FASL\_OBJECT and variables of type FASL\_OBJECT may be initialized, in which case that constructor will be called. Note also that the member variables of FASL are not directly accessible by script code (declared as private).

#### Socket Object

This is a structure declared in the meta-scope but there are no instances declared in the meta-scope as there is with FASL\_OBJECT. It is mostly like a normal FASL structure, but the constructor creates a windows TCP or UDP socket depending on what you pass in for the first argument of the constructor and the destructor cleans up that object. The member functions permits data to be sent or received on that socket, assuming it connected properly. The structure is as follows (the data members are omitted as it is not necessary to know them):

```

structure Socket
{
    function Socket(PROTOCOL eProtocol, int nPort)
    {
        // eProtocol is either TCP or UDP and the port is
        // the IP port.
        // Creates a windows socket object.
    }
    function ~Socket( )
    {
        // Cleans up the Windows socket object associated
        // with this socket.
    }
    function boolean socketOpen( )
    {
        // Typically must be called before sending or
        // receiving data for either protocol.
        // if returns false then it could not open the
        // socket and communication
        // will fail.
    }
    function boolean socketClose( )
    {
        // Close the socket. destructor will do this if
        // not done here.
    }
    function void BindPort(int nPort)
    {
        // Use when you want specify the source port
        // explicitly.
        BIND_PORT(nPort);
    }
    function string socketRead(int nBytes, int nTimeout)
    {
        // Read nBytes from Socket... returns empty string
        // if fail.
    }
    function string socketReadLine(int nBytes, int
    nTimeout)
    {
        // Read nBytes or up to end of line whichever
        // comes first.
        // return empty string on fail.
    }
    function binary socketReadBinary(int nBytes, int
    nTimeout)
    {
        // Read nBytes of raw binary data. empty binary on
        // fail.
    }
    function int socketWrite(string szBuffer, int nLength)
    {
        // Write nLength characters out of szBuffer.
        // return 0 on fail
        // or number of bytes written on success.
    }
    function int socketWrite(string szBuffer)
    {
    }
}

```

-continued

```

        // Write the entire string (length implicit),
        otherwise same as above.
    }
    function int socketWrite(binary bBuffer)
    {
        // Write the binary raw data... length implicit,
        otherwise same as above.
    }
};

```

#### Vulnerability Object for Active Assessment

This object is used for active assessment, and its custom behavior—atypical of general FASL script—is that it knows how to copy itself, and all the data it needs to find itself later, into a global data area where other scripts can access the data, and use the information to attempt to compromise another machine. The object is as follows:

```

structure Vulnerability
{
private:
    int m_nFaultlineID,
        m_nIPAddress,
        m_nExploitIPAddress;
    string m_szDescription;
public:
    function Vulnerability( )
    {
        m_nFaultlineID = 0;
        m_nIPAddress = 0xFFFFFFFF;
        m_szDescription = "Vulnerability
uninitialized.";
    }
    function void addToExploitableData(
        int nFaultlineID,
        int nIPAddress,
        int nExploitIPAddress,
        string szDescription)
    {
        // This sets all the member variables of this
        structure. This function may be
        // called from a derived structure and this
        function will know that. It stores the
        // entire contents of the object as well as the
        object's typename and size in the
        // global vulnerability data area.
    }
    function boolean getExploitableData(int nIndex)
    {
        // this function searches the global
        vulnerability
        // data area for the nIndex-the instance of a
        variable of the same type as this (this
        // could be and probably will be a derived
        structure from Vulnerability) and copies
        // its contents into this object. If there is no
        nIndex-th object, return false and
        // no change to this object's data will have
        occurred.
    }
    // Accessor functions, the members of this structure
    need to be read only once they
    // have been stored with addToExploitableData( ).
    function int getFaultlineID( )
    {
        return m_nFaultlineID;
    }
    function int getIPAddress( )
    {
        return m_nIPAddress;
    }
    function int getExploitIPAddress( )
    {

```

-continued

```

        return m_nExploitIPAddress;
    }
    function string getDescription( )
    {
        return m_szDescription;
    }
};

```

Behind this object is the concept that one script may discover something that another script can use to compromise another machine. Any vulnerability needs the information contained in this base class, in this embodiment. The m\_nFaultlineID is the vulnerability id of the vulnerability discovered. The m\_nIPAddress is the machine it was discovered on. The m\_nExploitIPAddress is the machine on which the data was discovered that proved instrumental in finding this vulnerability. The m\_szDescription is what you want to appear in the Active-Assessment report. The m\_nExploitIPAddress should typically be set to -1 when no other vulnerability was used to find this vulnerability. When another vulnerability is used, its m\_nIPAddress (typically using othervuln.getIPAddress( )) is entered into this vulnerability's m\_nExploitIPAddress. This will establish an audit trail which the report can represent graphically.

The general, usage of this is to derive a structure from Vulnerability, (i.e., "structure UnicodeVulnerability { // extra info specific to unicode vuln that can be used by other scripts};"). When a unicode vulnerability is found, for one example, a variable such as "UnicodeVulnerability uv;" is created, and its extra data is set and added via a call to "uv.AddToExploitableData( . . . )." After this call, another script that looks to make use of this particular vulnerability has code resembling the following:

```

int i;
UnicodeVulnerability uv;
for (i = 0; uv.getExploitableData(i); i++)
{
    // Attempt to use this data to compromise this host.
    if (succeed at compromising this host)
    {
        // create a vulnerability representing this host's
        vulnerability that was
        // found using the UnicodeVulnerability. Note
        that this vulnerability
        // may or may not be a UnicodeVulnerability... it
        could be some other
        // vulnerability. When you have populated it with
        its specific data, call:
        newvuln.addToExploitableData(
            nNewVulnFaultlineID,
            FASL.getAddress( ),
            uv.getAddress( ),
            "we got this vuln by exploiting Unicode on
another machine");
    }
}

```

#### Debug messages

There is a function called "debugMessage(<scalar>e)" overloaded for all the scalar types. Mostly it prints what is expected from a debugging system. Int and char print as numbers, boolean prints as "true" or "false", strings print their contents. Binary produces a hex dump output, through debugmessage(binary b), which closely resembles what the MS-DOS DEBUG.EXE "d" option produces. That is, it outputs out lines representing 16 bytes each and the line

format is "<4 digit hex offset>:<8 2-digit hex numbers separated by space>--<the other 8><16 chars where printable chars are represented as is and non-printable are represented as periods>". For example:

"0060: 67 65 6C 73 20 6C 6F 6F-6B 20 6C 69 6B 65 20 74  
gels look like t"

With the functions about to be specified, it is easy to convert other types into strings so you can `debugMessage` entire expressions . . . i.e. `debugMessage("here is what x equals:"+intToString(hex))`. There is typically no need to remove or comment out `debugMessage` calls in scripts that get "signed off," as final as the system execution of scripts will typically ignore the debug output.

## String functions

```

function string leftTrim(string s)
{
    // Lops off whitespace up to the first non-whitespace
    character.
}
function string rightTrim(string s)
{
    // T.ops off all trailing whitespace.
}
function int stringLength(string s)
{
    // self-explanatory.
}
function string toLower(string s)
{
    // Makes any characters that are upper-case lower-case.
}
function string toUpper(string s)
{
    // Makes any characters that are lower-case upper-case.
}
function int stringToInt(string s)
{
    // similar to atoi in C.
}
function string intToString(int n)
{
    // i.e 1048576 becomes string "1048576"
}
function string intToString(char c)
{
    // Similar to intToString(int).
}
function string intToHexString(int n)
{
    // i.e. 16383 becomes "3FFF"
}
function int hexStringToInt(string s)
{
    // i.e. "3FFF-blahblablah" becomes 16383.
}
function string intToBinaryString(int n)
{
    // i.e. 85 becomes "1010101"
}
function int binaryStringToInt(string s)
{
    // i.e. "1010101blablablah" becomes 85.
}
function string intToIPString(int n)
{
    // i.e. 16777215 becomes "0.255.255.255"
}
function string grep(string s1, string s2)
{
    return "to be implemented";
}
function int locate(string s1, string s2)
{
    // return 0-based position of s1 in s2 or -1 if s1 not

```

-continued

```

in s2.
}
5 } function string subString(string s, int nStart, int
nNumChars)
{
// i.e. subString("one law for the lion and ox is
oppression", 4, 3) = "law"
// it is smart about boundaries... if you ask for more
10 characters than exist
// you only get the ones that exist.
}
function string garbageString(char c, int nLength)
{
// i.e. garbageString('A', 7) = "AAAAAAA"
15 }
function string garbageString(int nLength)
{
// return string of nLength length whose characters are
random
// upper and lower case alphanumerics.
20 }
}
-----
Binary functions:
-----
function int binaryLength(binary b)
{
// self-explanatory.
25 }
function char binaryNthByte(binary b, int n)
{
// return nth byte numbered from 0. If out of bounds,
// return -1.
}
30 function boolean binaryChangeByte(binary& b, int n, char c)
{
// Change byte #n in binary b to c. If out of bounds,
// do nothing to b and return false.
}
function binary binaryRight(binary b, int n)
35 {
// return a binary which is the rightmost n bytes of b.
If there
// are not n bytes in b, you get all of b (not padded
to n).
}
40 function binary binaryLeft(binary b, int n)
{
// same like binaryRight except from left.
}
function binary binaryMid(binary b, int nStart, int nLength)
{
// nStart is 0-based... return binary which is the
45 nLength bytes starting
// at nStart. If there are not nLength bytes, return
however many there are.
}
}
-----
General global functions:
50 -----
function string getLocalHostIP( )
{
// return string "xxx.xxx.xxx.xxx" representing the IP
of the machine
// on which the FASL script is running.
55 }
function string getTargetHostIP( )
{
// This is in here for comparibility reasons... you can
get the same
// effect by doing:
60 intToIPString (FASL.getAddress( ) );
}
function int getTargetPort( )
{
return FASL.m_nDefaultPort;
}
function boolean RPCCheck(int nRPCProgNum)
65 {
// Attempt to make an RPC call on the given nRPCProgNum

```

-continued

```

to see if it exists
// on the target machine. Return true if it does.
}
function boolean SMBCheck(string szUserName, string
szPassWord, string szShare)
{
// Attempt to do a "net use" sort of thing on the given
share... i.e. "IPC$",
// "\\hostname\ipc$", "\\xxx.xxx.xxx.xxx\ipc$", etc. If it
succeeds then
// promptly delete the share and return true, otherwise
return false.
}

```

This implementation of the FASL scripting language permits one to perform active assessment of potential vulnerabilities on a known target computer with known open ports. For example, a FASL script to test a unicode vulnerability has the form:

```

structure UnicodeVulnerability extends Vulnerability
{
string m_szUnicodeString;
};
function fastmain( )
{
UnicodeVulnerability uv;
uv.m_szUnicodeString = "The unicode string";
uv.addToExploitableData(1, FASL.m_nIPAddress, "the
string for reporting purposes");
debugMessage("vulnfind executed.");
}
structure UnicodeVulnerability extends Vulnerability
{
string m_szUnicodeString;
};
function fastmain( )
{
UnicodeVulnerability uv;
if (uv.getExploitableData(0) )
{
debugMessage("getExploitableData( )...
m_nFaultlineID = " + intToString(uv.m_nFaultlineID) + ",
m_nIPAddress = " + intToIPString(uv.m_nIPAddress) + ",
m_szDescription = " + uv.m_szDescription + ",
m_szUnicodeString = " + uv.m_szUnicodeString);
}
else
{
debugMessage("getExploitableData( ) failed.");
}
}
}

```

Notably, information collected from a successful vulnerability test, in one embodiment, is used to advantageously further test the vulnerability of the target network and individual target computers. In particular, in the event of a successful vulnerability test, additional ports or levels of account access are typically available on the target computer. Thus, even if other prior vulnerability tests failed, they are advantageously retried after a successful vulnerability test.

FIG. 8 illustrates one embodiment of active assessment of a vulnerability of a target computer on a target network. For simplicity, the network is shown as having a target 1 computer 800 and a target 2 computer 802. For simplicity, a single vulnerability is assumed to apply to all TCP and UDP ports of the two computers. A single vulnerability 65 "TEST" is tested on various TCP and UDP ports on both target computers. Prior host discovery and port discovery

have provided target 1 data 804 and target 2 data 806 that includes an identification of open ports found on each target computer. Specifically, in the example illustrated in FIG. 8, TCP port 80 and UDP ports 5000 and 53 were found open on the target 1 computer 800, and TCP ports 23 and 80 and UDP ports 500, 53, and 1721 were found open on the target 2 computer 802. The active assessment process starts in a step 806 and begins by executing a TEST routine 808 that applies the TEST vulnerability to each port on the target 1 computer 800. In the example of FIG. 8, in a first round 810 of testing the target 1 computer 800, the testing of ports 80, 5000 and 53 result in no successful applications of the vulnerability. The system then moves in to a first round 812 of testing the target 2 computer, where testing of ports 80 and 53 are unsuccessful, but testing of ports 23, 5000 and 1721 are successful. Upon successful testing, an attempt is made to strip useful information from the target computer, and any information retrieved is stored in the target 2 data 806. Based on the new information retrieved, a second round 814 of testing the target 1 computer is attempted. In this attempt, testing of port 80 is still unsuccessful, but testing of ports 5000 and 53 is now successful with the inclusion of the additional information gathered from the target 2 data 806. In this manner, a second round 816 of testing the target 2 computer is attempted. The rounds of testing are repeated until vulnerability assessments can be actively completed for the entire target network.

Known vulnerabilities are typically stored in a database for inclusion in vulnerability assessment testing, active assessment and reporting. In one embodiment, the "Vulns" table called "static data," which only changes when deploying new patches incorporating new vulnerabilities. This would happen, for example, as new vulnerabilities are discovered and existing in-the-field system installations need to be made aware of them. The present system uses, in one example, the VulnsFound table below to indicate which vulnerabilities it has found on a particular scan. VulnsFound is simply an index into the Vulns table which keeps us from otherwise repeating the same data a multitude of times and wasting space.

TABLE: Vulns

COLUMN: BID, TYPE: varchar(10)
COLUMN: CVE, TYPE: varchar(20)
// "Common Vulnerabilities and Exposures" number.
COLUMN: CyberID, TYPE: int(10)
COLUMN: Description, TYPE: text(2147483647)
COLUMN: ExploitDataType, TYPE: varchar(64)
// Name of the FASL type, where applicable, that contains extra data for this vulnerability.
COLUMN: ExploitDate, TYPE: smalldatetime(16)
COLUMN: ExploitLink, TYPE: varchar(255)
// Site where you can download the exploit/patch.
COLUMN: FaultlineID, TYPE: int(10)
// Primary key field... linked to VulnsFound table.
COLUMN: Impact, TYPE: tinyint(3)
COLUMN: ISSID, TYPE: int(10)
COLUMN: LHF, TYPE: bit(1)
// short term intrusion opportunity... true or false.
COLUMN: Name, TYPE: varchar(128)
COLUMN: NID, TYPE: int(10)
COLUMN: Observation, TYPE: text(2147483647)
COLUMN: Person, TYPE: varchar(50)
COLUMN: Popularity, TYPE: tinyint(3)
COLUMN: Recommendation, TYPE: text(2147483647)
COLUMN: Risk, TYPE: tinyint(3)
COLUMN: RiskText, TYPE: text(2147483647)
COLUMN: Simplicity, TYPE: tinyint(3)
COLUMN: Type, TYPE: varchar(50)

-continued

//web, router, unix, trojan, etc.

TABLE: VulnsFound

COLUMN: ConfigurationID, TYPE: int(10)  
 COLUMN: CustomerID, TYPE: int(10)  
 COLUMN: FaultlineID, TYPE: int(10)  
 COLUMN: HostID, TYPE: int(10)  
 COLUMN: JobID, TYPE: int(10)  
 COLUMN: VulnFoundID, TYPE: int identity(10)

## VII. Quantitative Scoring of Target Network Vulnerability

Network vulnerabilities have traditionally been rated on a "low risk," "medium risk," and "high risk" scale. This subjective scale is based on the level of access to a target system granted by a vulnerability, the ease of detecting and exploiting the vulnerability, the public knowledge of a vulnerability, the percentage of computers subject to the vulnerability, and the like. However, this subjective scale lacks objective indicia allowing comparison of otherwise heterogeneous networks for the purpose of comparing relative security. Further, this subjective three-level scale provides little information about security improvement or added risks over time.

In one embodiment, the present system provides objective, quantitative indicia of the overall security of the target network. Advantageously, this quantitative indicia can be any quantitative scale of sufficient granularity to provide meaningful distinctions due to changes in network vulnerabilities over time, with a range of between 1-10, 1-100, and the like. Further, this objective indicia applies a standard formula to the various vulnerabilities discovered on the target network and network configuration, such that a valid comparison can be made between the security of otherwise heterogeneous network configurations, operating systems, and computers.

For example, the objective indicia is a risk measurement algorithm, such as a FoundScore F. In one embodiment illustrated by the flowchart in FIG. 9, the FoundScore F is defined as:

$$F = 100 - V - E \quad (\text{Eq. 1})$$

where,

F=FoundScore

V=Vulnerability Loss

E=Exposure Loss.

More specifically, in one preferred embodiment, the vulnerability loss V is defined as the sum of values assigned to each vulnerability found on the target network. For each of the n vulnerabilities found on the target network, that vulnerability is assigned a weight.

Thus, for a particular vulnerability  $V_x$ , where  $1 \leq x \leq n$ , the vulnerability weight  $V_{w,x}$  is defined as:

$$V_{w,x} = \begin{cases} \{50 \mid PEA(V_x) \geq 2\} \\ \{10 \mid 2 > PEA(V_x) \geq 1\} \\ \{5 \mid 1 > PEA(V_x)\} \end{cases} \quad (\text{Eq. 2})$$

where:

$$PEA(V_x) = \text{floor}((P(V_x) + E(V_x) + A(V_x))/3) \quad (\text{Eq. 3})$$

and:

floor(x) is the standard floor function,

 $P(V_x)$  is the popularity of the vulnerability on a 1-3 scale,

$E(V_x)$  is the ease of exploitation of the vulnerability on a 1-3 scale,

$A(V_x)$  is the level of access/privileges granted by the vulnerability on a 1-3 scale,

where the higher the score on the scale the greater the risk.

Alternatively, the vulnerability can be simply defined as the risk level associated with the vulnerability, such as where:

$$PEA(V_x) = R(V_x) \quad (\text{Eq. 3a})$$

where,

$R(V_x)$  is the risk associated with the vulnerability on a 1-3 scale.

Typically, the factors used to determine the  $PEA(V)$  for a particular vulnerability are provided from the vulnerability database discussed previously. In other embodiments, the vulnerability base function (e.g.,  $PEA(V)$ ) can be defined from a number of different variables, including, for example, ease of distribution of a vulnerability, ease of preventing the vulnerability, and the like. Thus, the total vulnerability score V is, in one embodiment, equal to:

$$V = \min(50, \sum_{x=1 \rightarrow n} \{V_{w,x}\}) \quad (\text{Eq. 4})$$

where:

n=number of vulnerabilities found on the target network,

Σ=the summation symbol,

 $V_{w,x}$ =the vulnerability weight defined above, and,

min(x,y)=the standard minimum function.

The total exposure score E is, in one embodiment, defined as:

$$E = \min(50, \sum_{y=1 \rightarrow q} \{10 \times U_y + 5 \times I_y + N_y + 2 \times M_y\}) \quad (\text{Eq. 5})$$

where:

q=the number of live target computers found on the target network,

$U_y$ =the open UDP ports found on the yth target computer, except DNS,

 $I_y$ =the open ICMP ports found on the yth target computer,

$N_y$ =the non-essential services found on the yth target computer,

$M_y$ =a penalty for target computers with no essential services,

Σ=the summation symbol, and,

min(x,y) is the standard minimum function.

In particular, the number of live target computers q is typically deduced from the host discovery process described above. The number of open UDP ports  $U_y$  found on the yth target computer is typically deduced from the service discovery and active assessment process described above. The number of non-essential services  $N_y$  found on the yth target computer is typically deduced from the service discovery of TCP and UDP ports, and in one preferred embodiment, counts all open ports found on each target computer, except for certain predetermined services. Table 11 lists examples of predetermined services that are not included in the count of nonessential services  $N_y$  in accordance with one embodiment of the present invention.

TABLE 11

Example Essential Services not Included in $N_y$ count	
Essential Service	Ports
HTTP	80, 8000, 8080
SSL/HTTPS	443
FTP	21
DNS	53
SMTP	25

Other combinations of essential services are also possible, based on changes in operating system, protocols, and the like. Finally, the no essential services penalty flag  $M_y$  is set to 1 for each target computer that has  $N_y > 0$ , and has no essential services running such as those shown in Table 7. Otherwise,  $M_y$  is set to 0.

As a result, the FoundScore  $F = 100 - V - E$  provides a quantitative score between 0 and 100 representing an objective indicia of the relative security of the target network. Generally, the higher the Foundscore (i.e., the higher  $F$  is), the greater the relative security of the target network. Other embodiments are possible, such as, for example, where:

$$E_{alt} = [\sum_{i=1 \rightarrow q} \{(5 \times U_i + 2 \times I_i + 2 \times N_y \times 2 \times M_y)^2\}]^{1/2}/q \quad (\text{Eq. 7})$$

$$V_{alt} = [\sum_{i=1 \rightarrow n} \{(V_{w,x})^2\}]^{1/2}/n \quad (\text{Eq. 8})$$

$$F_{alt} = E_{alt} + V_{alt} \quad (\text{Eq. 9})$$

In this alternative scoring regiment, the higher the  $F_{alt}$  score the worse the target network security.

FIG. 9 illustrates one embodiment of a methodology for determining the security score for a target network. In a first decision step 900, the method determines whether all vulnerabilities found in the target network have been counted. If not, the method calculates  $PEA(V_x)$  (Eq. 3) above, or a variation thereof, in a step 902 for vulnerability number  $x$  based on predetermined values stored in a vulnerability database 904. The value of  $PEA(V_x)$  is used to calculate the weighted vulnerability  $V_{w,x}$  (Eq. 2) for a vulnerability  $x$  in a step 906. Typically, the vulnerability counter is then incremented in a step 908 if necessary for a particular embodiment. Thereafter, the method returns to the decision step 900 and again determines if all the vulnerabilities have been counted.

After all the vulnerabilities have been counted, then in preferred embodiments, the total vulnerability  $V$  is calculated in a step 910. As discussed above, the total vulnerability  $V$  is the lesser of either 50 or the sum of all weighted vulnerability scores  $V_{w,x}$  in this embodiment. Alternatively, other scores are possible, or the tabulation of the total score  $V$  can be combined with the prior loop.

Then, in a decision step 912, the method determines whether all target computers found in the target network have been counted. If not, the exposure values are determined in a step 914. In preferred embodiments, the exposure values include  $U_y$  (the open UDP ports found on the  $y$ th target computer, except DNS on port 53),  $I_y$  (the open ICMP ports found on the  $y$ th target computer), and  $N_y$  (the non-essential services found on the  $y$ th target computer), are determined. The exposure values are determined dynamically based on the network security test, and are, in one embodiment, stored in the target computer vulnerability database 904 or in another database. In a step 916, a penalty  $M_y$  is determined for target computers with no essential services present. Preferably, the exposure counter is then

incremented in a step 918, and the method returns to the decision step 912 to determine whether all the target computers have been counted.

When all the target computers are counted, the total exposure  $E$  is calculated in a step 920 as the lesser of either 50 or the sum of a weighted sum of the exposure values found. Then, in preferred embodiments, the score  $F$  is calculated in a step 992 as 100 minus the total vulnerability and exposure scores to generate a representation of a network security score. In preferred embodiments, a greater value of  $F$  implies greater network security on an objective scale.

Other exemplary embodiments apply various permutations of the factors associated with vulnerability scoring and exposure scoring, and are similarly foreseen as within the disclosure herein.

#### VIII. Graphical Hierarchical Reporting of Target Network Topology and Vulnerabilities

Generally, a conventional network security system provides reporting information in a text dump format. The collected data records regarding network topology, target computers, vulnerabilities, and the like are dumped into an ASCII file that requires substantial effort to interpret. Such data dumps are used conventionally because, other than alerting the network security user to the presence of high danger vulnerabilities, existing systems do not provide an interactive, hierarchical, and graphical representation of the data representing the network, target computers, and vulnerabilities found to assist the user in identifying and correcting the specific vulnerabilities.

The present system compiles the data discovered during security testing into a graphical, informationally hierarchical, and interactive set of documents for review at various levels of detail and documentation. Thus, in one embodiment of the present invention, the reports engine produces (1) a dynamic graphical display representing the network topology, the target computers found, and the vulnerabilities found throughout the target network; (2) a comprehensive list of target computers, vulnerabilities found, and vulnerability explanations; (3) an objective scoring report describing the approximate security rating of the target network; (4) an exposure report describing the ports, services, and (5) detailed information that sets forth testing results on a per-machine, per-port, or per-vulnerability basis. Certain individuals within an organization may want different levels of detail. For example, upper management may only want the objective scoring report (3) that describes the approximate security rating of the target network. In contrast, the network administrator wants to receive all reports, particularly, the detailed information report (5) that enables the administrator to identify the machines and ports on the machines that need to be corrected.

For example, a functional representation of a report generated by the report engine typically includes a quantitative score representation including (1) the actual score, (2) a color representation of the score exposure, and (3) a graphical representation of the quantitative score. The actual score is the  $F$  score or an alternative embodiment, such as, for example, a 1 to 100 numerical rating as described previously. The color representation represents the overall vulnerability range of the score. Thus, in one embodiment, a score between 1 and 33 will have a red color signifying high vulnerability, a score between 34 and 66 will have a yellow color signifying medium vulnerability, and a score above 66 will have a green color signifying low vulnerabil-

ity. Other colors, ranges, and representations, such as, for example, using icons or pictures to represent the vulnerability level, are foreseen.

FIG. 10, for example, illustrates an embodiment of a hierarchical security report 1000 presented on a display for viewing by a user. The hierarchical security report includes a representation 1010 of the Foundstone score, showing a value of 15 corresponding to a relatively high risk. The security report further includes a graphical representation 1020 of the mapped network topology, including the live target computers and the services located on the target computers. The security report further includes a list 1030 of discovered hosts, and a list 1040 of discovered services.

The user may select any of the reports shown in FIG. 10 (or other reports that are not shown) and request additional details. FIG. 11, for example, illustrates an embodiment of a hierarchical security report 1100 that provides greater detail regarding the vulnerabilities of two target computers. For example, an upper portion of the display identifies a first target computer ("localhost.local.net" at IP address 255.255.255.255) having a first open port represented by a window 1110 and a second open port represented by a window 1120. Each of the open ports of the first target computer has a respective vulnerability and may identify a vulnerability patch to install on the target computer to reduce the vulnerability. As a further example, a lower portion of the security report 1100 identifies a second target computer ("localhost2.local.net" at IP address 254.255.255.255) having a single open port represented by a window 1130. The information in the window 1130 indicates that the service on the port is out of date, suggesting that it should be removed, and advantageously identifies a vulnerability patch to install if the service is to be retained.

There is an "object layer" implemented in C++ which represents the hierarchy of data objects that the system deals with, as described above with respect to the FASL scripting language. All objects follow the model of 1) containing the data directly relevant to that object as data members (these are 1-1 correspondent with a row in the database) and 2) possibly containing lists of other objects that also follow this convention. Each object has a Load() and Save() member that deals with the data for which that object is directly responsible by loading or saving it from the database. For the "child" object lists, the UpdateToDatabase() and UpdateFromDatabase() member functions are used. These call the object's Save() or Load() members respectively. They then recursively call the child lists' UpdateToDatabase() or UpdateFromDatabase() members. In this way, objects and their children can be selectively loaded from any point in the hierarchy. This was done because the objects are generally too unwieldy to deal with all at once.

Once the data is loaded, an object which represents a single scan has a GenerateReport member function which programmatically generates all the html, gif, jpg, and png files that constitute the report. The html files are generated by loading a "template" html file that contains boilerplate source and key words and replacing the key words with the html source representing the report data. The gifs and jpgs are simply background graphics that do not change from report to report so they are written as is without the report mechanism having to know the details of gif and jpg encoding. For images that do need to be dynamically calculated, the png graphics format is used. This is a public (non-licensed) format for which there exists a public library which we are linking into the system executable. The images are drawn onto a Windows HDC (windows software object

that represents a display device) and then converted into png by a custom class which wraps the object in the proper format.

An exemplary *Scan Summary Report* is set forth below in Appendix A. A corresponding exemplary *FoundScore Report* is set forth below in Appendix B. A corresponding *Network Topology Report* is set forth below in Appendix C.

Although the above description of the present invention includes at least one embodiment of the system, it should be understood that various changes, substitutions and alterations can be made to the system without departing from the spirit and scope of the invention as defined by the claims herein. For example, one alternative embodiment of the methodology described above is exemplified in FIG. 12.

What is claimed is:

1. A system for determining an operating system of a target computer operably connected to a network, the system comprising:

first and second data packets, said first and second data packets compliant with a protocol supported by said network, said first and second data packets transmitted via said network to said target computer;

first and second operating system fingerprints comprising data bits stored in a computer-readable medium, said first and second operating system fingerprints associated with a first operating system;

a first target computer fingerprint comprising data bits stored in a computer-readable medium, said first target computer fingerprint including a representation of at least a portion of data received in response to said transmission of said first data packet;

a second target computer fingerprint comprising data bits stored in a computer-readable medium, said second target computer fingerprint including a representation of at least a portion of data received in response to said transmission of said second data packet; and

fingerprint comparison instructions embodied in a computer readable storage medium, said instructions executable by a computer to compare said first operating system fingerprint and said first target computer fingerprint, to compare said second operating system fingerprint and said second target computer fingerprint, and to generate a result indicative of whether said first operating system was running on said target computer; wherein the first and second data packets each include TCP packets.

2. The system as described in claim 1, wherein a first range of bits of said first data packet represents a first parameter value, and wherein said first range of bits of said second data packet represents a second parameter value different from said first parameter value.

3. The system as described in claim 2, wherein said second parameter value is derived by changing one bit in said first range of bits of said first data packet.

4. The system as described in claim 2, wherein said first and second operating system fingerprints differ.

5. The system as described in claim 4, further comprising: a third data packet, said third data packet compliant with said protocol, said first range of bits of said third data packet representing a third parameter value different from said first and second parameter values, said third data packet transmitted via said network to said target computer;

a third operating system fingerprint comprising data bits stored in a computer-readable medium, said third operating system fingerprint associated with said first oper-



61

ating system, said third operating system fingerprint differing from said first and second operating system fingerprints; and

a third target computer fingerprint comprising data bits stored in a computer-readable medium, said third target computer fingerprint including a representation of at least a portion of data received in response to said transmission of said first data packet, said comparison instructions executable by a computer to compare said third operating system fingerprint and said third target computer fingerprint before generating said result.

6. The system as described in claim 5, further comprising: fourth, fifth and sixth operating system fingerprints comprising data bits stored in a computer-readable medium, said fourth, fifth and sixth operating system fingerprints associated with a second operating system, at least one of said fourth, fifth and sixth operating system fingerprints differing from a respective one of said first, second and third operating system fingerprints;

said comparison instructions executable by a computer to compare said fourth operating system fingerprint and said first target computer fingerprint, to compare said fifth operating system fingerprint and said second target computer fingerprint, to compare said sixth operating system fingerprint and said third target computer fingerprint, and to generate a second result indicative of whether said second operating system was running on said target computer.

7. The system as described in claim 5, wherein said protocol is TCP/IP and wherein said first range of bits corresponds to a packet field representing a maximum segment size.

8. The system as described in claim 5, wherein said first parameter value is obtained by setting no bits, said second parameter value is obtained by setting one bit, and said third parameter value is obtained by setting two bits.

9. The system as described in claim 5, wherein said first parameter value is 0, said second parameter value is 128, and said third parameter value is 128 plus a multiple of 256.

10. The system as described in claim 5, wherein said first range of bits represents at least two bytes, and wherein a value of said second parameter is obtained by setting the last bit in a byte, and a value for said third parameter is obtained by setting the last bit in a byte.

11. The system as described in claim 10, wherein said third parameter is obtained by setting adjacent bits in said first range of bits.

12. The system as described in claim 5, wherein said first, second and third data packets are transmitted in order of lowest parameter value first.

13. A method for identifying an operating system of a target computer via a network, the method comprising the steps of:

sending a first data packet to said target computer via said network, said first data packet complying with a protocol of said network and having a first pattern of bits in a first range of bits;

generating a first response value representing at least a portion of data received via said network in response to said sending of said first data packet;

sending a second data packet to said target computer via said network, said second data packet complying with said protocol and having a second pattern of bits in a first range of bits, said second pattern of bits different from said first pattern;

62

generating a second response value representing at least a portion of data received via said network in response to said sending of said second data packet;

sending a third data packet to said target computer via said network, said third data packet complying with said protocol and having a third pattern of bits in a first range of bits, said third pattern of bits different from said first or said second pattern;

generating a third response value representing at least a portion of data received via said network in response to said sending of said third data packet;

comparing said first response value to a first predetermined value associated with a first operating system;

comparing said second response value to a second predetermined value associated with said first operating system;

comparing said third response value to a third predetermined value associated with said first operating system; and

generating a value indicative of a relationship between said first operating system and said target computer; wherein the first, second, and third data packets each include TCP packets.

14. The method as described in claim 13, the method comprising the further steps of:

comparing said first response value to a fourth predetermined value associated with a second operating system;

comparing said second response value to a fifth predetermined value associated with said second operating system; and

comparing said third response value to a sixth predetermined value associated with said second operating system.

15. The method as described in claim 13, wherein no bit is set in said first pattern of bits, wherein one bit is set in said second pattern of bits, and wherein two bits are set in said third pattern of bits.

16. The method as described in claim 13, wherein the number of bytes in said second pattern of bits that have at least one bit set is greater than the number of bytes in said first pattern of bits that have at least one bit set, and wherein the number of bytes in said third pattern of bits that have at least one bit set is greater than the number of bytes in said second pattern of bits that have at least one bit set.

17. The method as described in claim 13, wherein no byte in said first pattern of bits has a least significant bit or a most significant bit that is set wherein at least one byte in said second pattern of bits has a most significant bit that is set, and wherein at least one byte in said third pattern of bits has a least significant bit that is set.

18. The system as described in claim 5, wherein the third data packet includes an RFC-compliant TCP packet.

19. The system as described in claim 1, wherein the first data packet includes a TCP SYN packet with a maximum segment size MSS option in an options field thereof set to 0.

20. The system as described in claim 1, wherein the first data packet includes a TCP SYN packet with a maximum segment size MSS option in an options field thereof set to 128.

\* \* \* \* \*

**Appendix C**

***Evidence Appendix***

Other than the reference attached to the Appeal Brief as Appendix B, no evidence was submitted pursuant to 37 C.F.R. §§ 1.130, 1.131, or 1.132, and no other evidence was entered by the Examiner and relied upon by Appellant in the Appeal.

**Appendix D**

***Related Proceedings Appendix***

As stated on page 3 of this Appeal Brief, to the knowledge of Appellant' Counsel, there are no known appeals, interferences, or judicial proceedings that will directly affect or be directly affected by or have a bearing on the Board's decision regarding this Appeal.